Market Failure: When do markets get you efficiency?

1. You Need Perfect Competition
    a. price-taking; large number of individuals and firms; equal power for all players conditional on their output (money).
        i. monopolists are not price takers, and produce too little to hit Utility Possibility Frontier.
2. You Need Complete Markets
    a. you can buy and sell anything, eg contingent commodities and insurance
        i. nobody will sell me insurance for journals rejecting my papers due to information problems (below)
3. You Need No Externalities
    a. action taken by an agent affects other agents, but not through prices
        i. I drive to work without taking into account the crap coming from my tailpipe that is then breathed by Fraser Valley residents.
4. You Need No Public Goods
    a. public goods are can be *shared* without cost, and cannot be consumed *exclusively*, and no agents can choose not to consume.
        i. clean air can be shared nearly freely, cannot be made exclusive and cannot be denied (except in *Total Recall*).
5. You Need No Increasing Returns to scale.
    a. If returns to scale are increasing for a firm, competitive firms don't know how much to produce; the industry will be monopolised by the biggest producer
6. You Need Perfect Information
    a. everybody knows everything about commodities and prices:
        i. perfect knowledge of quality (health, education?)
        ii. perfect knowledge of prices (all of them, for all qualities)
        iii. perfect knowledge of qualities and prices in the future. (savings).
    b. Big Problems in *Insurance* markets. eg, firms don't know who is sickly and who's healthy, so they don't know what price to charge for insurance policies to people. They charge a fixed price independent of your health (which they can't observe), so healthy people don't buy any because it seems too expensive. used cars.
    c. *Adverse Selection* is when private information makes the agents who want to transact exactly the agents that others are trying to avoid. Eg, people who are going to live a long time want annuities (contracts which pay a fixed amount till you die).
        i. private seller gets the shaft because they can't identify low risk people.
        ii. public seller uses monopoly power and political power to solve this problem: force everybody to buy a single insurance policy.
    d. *Moral Hazard* is when agents behave differently with insurance than without it. Eg, if I have employment insurance, I might not mind getting laid off.
        i. private seller gets the shaft because insured people act in ways that increase the cost of insuring them.
        ii. public seller uses political power to solve this problem: punish people for breaking rules (eg, toss rule-breakers off welfare).
    e. lots of public programs are in insurance markets:
        i. public pensions: insurance against outliving your savings

        ii. public health: insurance against expensive illness
        iii. employment insurance: insurance against job loss

7. You need an adequate market distribution of income
    a. Lots of people think that government's *only* function is to redistribute. This is false. Governments try to solve the above problems (they are constrained in their efforts (see Government Failure below). Sometimes in solving the above problems, they get nice redistributive effects, too. This could be intentional, or it could be a side effect. Eg, public pensions are highly redistributive towards the current elderly. Since they group up poorer than us (on average), this is progressive.
8. Further Reading: Nick Barr <u>The Economics of the Welfare State</u>.
9. Government Failure: What makes government production and financing difficult?
    a. pentagon is big wastrel. *secrets* lend themselves to closed bidding, and nepotism.
    b. governments often give hidden subsidies to specific firms or industries.
    c. Canadian public rails are no more expensive than private rails. Public hospitals are cheaper than private ones (quality might be different). Public insurance administrative costs are *way* lower than private insurer costs.
    d. Nice little book by Charles Wolf, Jr, <u>Markets or Governments</u> says that really we are not choosing between perfect markets or perfect governments, but rather between imperfect markets and imperfect governments.
        i. cost and revenue not connected. no disciplinary device like bankrupcty.
        ii. Individuals in organizations have their own (non profit maximising) goals.
        iii. sometimes programmes have unintended effects that the creating agency *is not responsible for*. Similar to firm making pollution
    e. **These problems exist in large organizations, whether public or private.**

Chapter 2

p 14 The marginal rate of time preference (MTPR) is the marginal rate of substitution (MRS) across periods.
The MTPR exists without a market. It is just one's willingness to pay (in present consumption) for future consumption.
If 1+r is the MTPR, then we can use r to discount future returns to their present value. If we discount only net returns, then we get a net present value (NPV). Define C_t as the net return in period t=0,...,T (T could be infinity). Net returns are positive if benefits exceed costs in period t, and (weakly) negative otherwise.
NPV is given by $\sum_{t=0}^{T} \frac{C_t}{(1+r)^t}$, or, in continuous time, $\int_{t=0}^{T} \frac{C(t)}{(1+r)^t} \P t$.
If the discount rate (1+r) is not identical from period to period, you could discount by r(t).
A good trick is that as T goes to infinity, if C is constant over time, $\sum_{t=1}^{\infty} \frac{C}{(1+r)^t} = \frac{C}{r}$

para 2: What if there's no credit market and c1=\$10 and c2=\$20? Is it still a good project? What is the MTPR here?

P 16 Perfect credit markets: If the interest rate i gives the price of period 2 consumption (for both borrowers and savers), then for everyone (who maximizes) MRS=MTPR=1+i.
This is because every individual has an incentive to adjust her consumption so that her MTPR gets closer to i if her MTPR does not equal i.

p 18 This book was written before Thatcher closed the pits (coal mines).

p 20 Are the two project evaluation strategies really identical? What if both projects have negative net present value (NPV)?

P 20 The internal rate of return (IRR) is the discount rate such that NPV=0.
IRR is unique and decisive for a single project only if all of its costs precede all of its benefits. 'Decisive' means that the IRR > i, the project is good.
If all costs do not precede all benefits, then it may be possible that a project has two solutions for the IRR.

p 22 Why is IRR1>IRR2 not decisive for 2 projects? Because an IRR of 20% on a \$10 investment is worse than an IRR of 10% on a \$100 investment.

p 23 Interdependent projects are easy. Consider projects A and B. You have four alternatives to consider: (1) A only; (2) B only; (3) A and B; and (4) neither A nor B.

Chapter 3 Opportunity Costs vs Accounting Costs

p 30 Make all the possible courses of action explicit. ‘Opportunity cost’ is what is foregone as a result of engaging in a course of action.
‘Opportunity costs are costs that are incurred if you do a project that are not incurred if you do not do the project.’

p 30 Depreciation is not opportunity cost. Depreciations sum to the purchase price; opportunity costs need not. A capital asset might have a large depreciation during a year, but might be completely useless to the firm and to others in alternative uses, thus having zero opportunity cost.

p 31 Train example
Costs: wages and fuel—accounting costs are exactly opportunity costs (unless you can’t get rid of your labour or fuel!).
Trains—if you look at the accounts, you get some depreciation allowance, which is not equal to what you get if you don’t do the project—the sale value of the train.
Benefits: tickets.

P 32 Marginal cost (MC) doesn’t generally equal average cost (AC), but MC is what matters when trying to determine the opportunity cost of changes in input use. How do you get information on MC? After all, firms don’t just go around changing scale for the benefit of the analyst. Presumably, if they did change scale a lot, it is because technology is changing, and so even the changes in scale can’t help you identify MC.
You could use input demand analysis if you have data on many firms (but you still wonder about identification).
You could break accounting component costs into fixed and variable costs, and then assume linear costs (cost=fixed+q*variable). This strategy is very common.

P 37 Price scaling for inflation is equivalent to increasing the nominal MRTP. Whatever you do, you should use all nominal or all real quantities in appraisal. If you are using reals, make sure goods prices are real.

sunk costs don’t matter.

Chapter 4 What Discount Rate?

p 43 With perfect capital markets, we don't need to distinguish between income and consumption because we can allocate our consumption across time periods as we see fit as long as we have the lifetime income to cover it.

P 44-5 People have different MRTPs. Which should you use? Should you aggregate them? If they all have access to a perfect capital market, then they will adjust their consumption profiles over time to equate their MRTPs with the interest rate. The interest rate is the price of holding money today instead of saving it for tomorrow. Thus, everyone has the same MRTP in competitive market equilibrium, and it is equal to the interest rate i.

P 44 Why don't we need to worry about the distinction between income and consumption? Because people put their consumption where they want subject to an income constraint if capital markets are perfectly competitive.

P 46 Unfortunately, capital markets are not perfect (any student loan holders among you?). Thus, people will disagree about NPVs, and disagree about which projects are best.

P 47 They will disagree over other arrangements, for example, about the optimal timing of benefits. Firms also have different borrowing and lending rates.
We have already seen that the exact choice of discount rate matters a lot; now we see that people may not agree on which one to use.

P 48 For current lenders, use the rate they lend at. For current borrowers, use the rate they borrow at. That is, use the rate they face at the margin.

P 48 Public agencies are a little different (even if they are trying to maximize profit—eg, commercial crown corporations like BC Hydro). They are often capital constrained—that is, they are not allowed to borrow as much as they like. Thus, they may resort to alternative financing arrangements that are similar, but not identical, to borrowing. For example, many public agencies lease their capital, so that the debt doesn't show up.

P 48 Sometimes governments explicitly state a discount rate to be used for project appraisal. Eg, in BC, the government mandated discount rate was for many years 10%.

The discount rate chosen can really matter. Eg, what is the difference in PV for a constant stream K at an interest rate of 2% vs at an interest rate of 4%?

Chapter 5: Uncertainty

p 54-5 Probabilities are necessarily subjective—some person supplies them.

P 56 First order stochastic dominance(FOSD) is satisfied if one cumulative distribution function (CDF) is entirely above another one. If *f(x)* is the density of net benefits *x*,

FOSD if $\int_{-\infty}^{T} f(x_1)\P x > \int_{-\infty}^{T} f(x_2)\P x \; \forall T$

If FOSD is satisfied, then the project with the higher CDF is preferred no matter what risk preferences are used in evaluation. FOSD only ranks some distributions, not all of them.

P 57 You might also rank by expected value (EV).

EV is equal to $\int_{-\infty}^{\infty} f(x)x\P x$, so EV dominance requires $\int_{-\infty}^{\infty} f(x_1)x_1\P x > \int_{-\infty}^{\infty} f(x_2)x_2\P x$.

Consider two projects that have the same EV, but one has *x* spread wider than the other.

The pdfs are centered at the same place, but one has fatter tails and a lower middle than the other.

If you draw the CDFs, you will see that the fat-tailed one is worse by FOSD.

P 58 The certainty equivalent (CE) of a project is the net benefit that, if received for sure, would result in the same utility (or shareholder value) as *f(x)*. Defining *u(x)* as utility, the certainty equivalent is given by $u^{-1}\left(\int_{-\infty}^{\infty} f(x)u(x)\P x\right)$.

The book says that the "certainty equivalent is approximately equal to the expected value ... [if] the gamble is small ...[and] other risks faced are not correlated". Are these conditions likely to be true? Does this just mean that the certainty equivalent converges to the expected value as the project becomes riskless (the gamble becomes small)?

P 60 Can risk be summarised in a risk premium? If the risk premium represents the probability of failure (of zero return), then it is just the multiplier to get expected NPV. Consider a consider probability of permanent failure *p*. Survival to period *t* occurs with probability $(1-p)^t$. For small *p* and *i*, this is close to discounting by *(1+i+p)* instead of *(1+i)*. If the risk premium represents uncertainty, then it is wrong, because risky benefits are worth less and risky costs cost more. Risk premia count risky costs as *less* costly.

You could use CE's in each period rather than net benefits to solve this problem. Eg, CE's of 50-50 gambles for net benefits in each period gives numbers that are less than EV's.

p 62 How much of the future should you use? A nice trick for cutting of infinite streams of benefits is to notice that, for constant benefit rates, a finite stream is equal to an infinite stream minus a discounted infinite stream. For example, the NPV of *x* from period 0 to 50 is $\sum_{t=1}^{50}\frac{x}{(1+i)^t} = \sum_{t=1}^{\infty}\frac{x}{(1+i)^t} - \frac{1}{(1+i)^{50}}\sum_{t=1}^{\infty}\frac{x}{(1+i)^t} = \left(1-\frac{1}{(1+i)^{50}}\right)\frac{x}{i}$.

Uncertainty, A Diversion

READ EATON, EATON and ALLEN, ch 17 for a treatment of uncertainty and risk

Definitions (all in discrete state space–no integrals)

Expected Value (EV): $EV = \sum_{s=1}^{S} p_s x_s$ where states are indexed $s=1,...,S$, state s occurs with probability p_s, and in state s, you get x_s.
The expected value is the amount you'd pay for a lottery defined by p_s and x_s if and only if you were *risk-neutral* (and your utility function was continuously twice-differentiable). Let us assume for the moment that x is a level of expenditure (or income).

Utility Function: $u(x)$ gives the utility level associated with the (single-valued) outcome x. If there is only one state, which happens with probability 1, then $u(x)$ gives the utility level for the outcome in this state.

Expected Utility Hyp (EUH): Utility over a lottery, U, is equal to the probability-weighted sum of utilities in each state

$$U(p_s, x_s) = \sum_{s=1}^{S} p_s u(x_s)$$

Risk-Neutrality: If the expected utility hypothesis is satisfied, then risk-neutrality holds if and only if u is linear, that is, if $u(x)=a+bx$ for some a,b. Risk-Neutrality implies that the MU of money is constant.

Certainty-Equivalent (CE): The CE is the amount that, if had for certain, would give the same utility as the lottery in question.

Cost Function: The cost function is the inverse of the utility function over x. That is, $C(u) = u^{-1}(x)$. Eg, if $u(x)=a+bx$, then $C(u)=(u-a)/b$.

CE given EUH: Given EUH, for any lottery, there exists an $\bar{x}$ such that

$u(\bar{x}) = U(p_s, x_s) = \sum_{s=1}^{S} p_s u(x_s)$. Solve for the CE, which is $\bar{x}$:

$$CE = C(\bar{x}) = C\left(U(p_s, x_s)\right) = u^{-1}\left(\sum_{s=1}^{S} p_s, u(x_s)\right).$$

Risk-Averse: Given EUH, if risk-averse, then the CE is less than the EV. (If CE=EV, then risk-neutral. If CE>EV, then risk-preferring.) Equivalently, given EUH, if u is strictly concave in x, then risk-averse; if u is linear in x, then risk-neutral; if u is convex in x, then risk-preferring.

Willingness-to-pay: The CE is one's willingness-to-pay for a lottery.

Arrow-Pratt Measure — Coefficient of Relative Risk Aversion $r = -\dfrac{\dfrac{\partial^2 u(x)}{\partial x^2}}{\dfrac{\partial u(x)}{\partial x}} x$.

This measures the speed at which the derivative of *u* changes. Risk aversion is about the distance between the CE and the EV. If the slope of *u* changes fast, then this distance is big; otherwise not.

CES utility function — constant-elasticity-of-substitution utility function is defined as

$$u^{CES} = x^q / q \quad q \neq 0$$
$$= \ln x \quad q = 0$$

, where the parameter $q \leq 1$.

The coefficient of relative risk aversion in this class is $r = 1 - q$, which is weakly greater than zero. This coefficient is big for utility that is very risk averse, and small for not very risk averse, and zero for risk-neutral.

CES cost function — *c(u)* for the CES utility function is $c(u) = u^{-1}(x) = (q u)^{1/q}$

CES Certainty Equivalent — Given the EUH and CES utility (with $q \neq 0$, the certainty equivalent $\tilde{x}$ satisfies

$$u(\tilde{x}) = \sum_{s=1}^{S} p_s u(x_s) = \sum_{s=1}^{S} p_s \frac{(x_s)^q}{q} = \frac{1}{q} \sum_{s=1}^{S} p_s (x_s)^q$$

$$\tilde{x} = \left(\sum_{s=1}^{S} p_s (x_s)^q \right)^{1/q}$$

Examples — Consider the lottery A with 1/3 probability of getting \$10 and 2/3 probability of getting \$50. The expected value is the probability weighted sum of these two outcomes, EV=\$36.67.

For a person with $q = -1$, $r = 2$ and

$$\tilde{x} = u^{-1}\left(\frac{1}{3} u(10) + \frac{2}{3} u(50) \right) = \left(\frac{1}{3}\frac{1}{10} + \frac{2}{3}\frac{1}{50} \right)^{-1} = \frac{150}{8} = 18.75$$

For a person with $q = 0$, $r = 1$ and

$$\tilde{x} = u^{-1}\left(\frac{1}{3} u(10) + \frac{2}{3} u(50) \right) = \exp\left(\frac{1}{3} \ln 10 + \frac{2}{3} \ln 50 \right) = 29.25$$

For a person with $q = 1$, $r = 0$ and $\tilde{x} = EV = 36.67$.

Importance — Cost-benefit analysis often uses the EV as the measure of a risky outcome, and then sticks this into the NPV calculation. However, if the CE is much less than the EV, this overvalues risky projects.

Chapter 6 Input Constraints

p 74 For any given input, eg, trucks, if you are free to buy as much as you want, the opportunity cost can be no greater than the price at which you can buy it. Further, if you have some trucks, and you can sell as many as you want, then the opportunity cost can be no less than the price at which you can sell them. If these prices are the same, then that price is the opportunity cost of the input.

- If the 'buy' and 'sell' prices are not equal, then the opportunity cost of the input is not so easily spotted. If you are forbidden from buying, then the upper bound is infinity; if you cannot sell, then the lower bound is zero. Not much help.

p 76 Maximising NPV implies maximizing NPV per unit of the constrained input.

- For example, if capital is unconstrained (you can borrow as much as you want), then you should do any and all projects with a risk-adjusted NPVs exceeding zero.
- But, if capital is constrained, then do (nonscaleable or limitedly scaleable) projects in order of their return *per dollar of capital*. This is easiest to see if you treat every dollar as separate. Use your first dollar on the best NPV projects and then go down the list.
- This requires prior choice of the discount rate. (If the timing of investment and return is different across projects, then changing the discount rate changes the ordering of NPVs).

P 77 The *shadow price* is the return to, or opportunity cost of, a rationed good. See Figure 6.1 and Table 6.2. In the figure, the step function is the demand curve and the vertical line is supply. The (nonmarket) shadow price is 29 cents, because that is the marginal value of that capital in its next best use (more of project B). This shadow price is not the market price of anything.

P 79 You could also draw on shadow discount rates, which would be the internal rates of return of each project. These are good for comparison to the actual interest rate faced. If capital is not constrained, then you want to do any project whose IRR exceeds *i*.

- If capital is constrained, then the IRR is not decisive. Projects have high IRRs if their benefits are front-loaded *or* if their benefits are large. Given the interest rate, a project may have a higher IRR while having a lower shadow price. For example, investing $1 this year for a $1.10 benefit next year has an IRR of 10%. Investing $1 this year for 10 cents next year and forever also has an IRR of 10%. However, if the interest rate is actually 5%, then the NPV of the former is about 5 cents, whereas for the latter, the NPV is 1$. So you want to do the latter project only if you only have one dollar.
- Why is there a whole chapter on capital rationing? Many public agencies are simply not allowed to borrow—hospitals in many Canadian jurisdictions, many US states are not allowed to borrow without special permission by referendum, and some aren't allowed to borrow at all.
- Why don't we use shadow discount rates? Because we maximize NPVs, not IRRs.

p 81 What if your capital is rationed in many (or all) periods? If rationing is of the same stringency in all periods, then you want the marginal shadow price of capital to be equal in all periods.

P 82 "A project with positive NPV (at taxpayers MTPR) should be done, no matter what the capital constraint." The reason is that it is 'self-financing'. If you could move all the returns to the present (through a bank, for example), then new capital 'magically' appears.

Chapter 7 The Objective in Cost-Benefit Analysis p 90 A potential pareto improvement (PPI) is possible if one can find a feasible allocation which yields a pareto improvement (PI) if resources are suitably (and costlessly) transferred post-facto. Consider a space of total income for two individuals. Draw an endowment E, and a wiggly line represent the set of feasible allocations. The feasible allocation which features the maximum *sum* of income is a PPI from *any* endowment in the feasible set, even if it is not a PI from that endowment.

Cost benefit analysis (in its standard form) seeks PPIs. As such, it typically seeks to maximize the sum of monetary (or monetized) rewards. Thus, cost-benefit analysis can be seen as indifferent to distribution.

The important thing here is that the post-facto costless transfers don't actually happen. If they did, and the allocations would be post-facto pareto improvements (not just potential). The test of whether or not an allocation A is a PPI compared to the endowment E is often called the 'compensation test'. This is because if the people in A can write cheques to each other in *compensation* for the change and still be better off than they were in E, A is a PPI over E.

P 91 I am hazy on Sugden and Williams' dichotomy between the two approaches to justifying the cost-benefit objective function.

P 91 The 'decision-making' approach allows that actual policy makers (presumably elected folks) need to know about the sum of monetary rewards to make their decisions, and will worry about other things (like distributional implications) when they make their decisions. Here, we should note that even strictly positive economists can make predictions about the distributional implications of policy choices. The policy maker can then use this information and her knowledge of what is best for society as a whole to make the decision.

P 92 Alternatively, the cost-benefit analyst might take what the authors call a 'paretian approach', wherein they try to inform the policy maker of the best project in a more universal sense. Rather than having the policy maker aggregate all the information (like NPVs and distributional and other consequences), the cost-benefit analyst plays the role of the policy maker and aggregates the considerations herself.

p 94 Welfare is comprised of efficiency and equity. "Cost-benefit analysis should look for the efficiency part only". However, if there are indifference curves over equity (an inequality index over money) and efficiency (total money), then this is equivalent to maximizing a welfarist social evaluation function over individual utility (which depends on money) (Donaldson 2002).

P 91,5 Even if you don't favour the potential pareto criterion due to its distributional indifference, you might argue that governments can make these potential compensations real by way of its coercive power. However, it seems to me that if you think the potential pareto improvements are actually feasible pareto improvements, then you should just call them that.

P 97 The perfectly competitive model zeroes out lots of effects, so it 'contains' the scope of welfare effects that need to be measured. For example, with competitive input markets, p=MC=Marginal social value.

P 98 Even if you think distribution matters, and you think potential compensations won't actually be done, there is still an argument to use the potential pareto improvement criterion. If governments already have the power to optimize the distribution of income and well-being, then they can redo the optimization after the project is done, and everything will be great again.

Pareto improvements make everyone weakly better off. There is general agreement that it is good to do this if possible. (I, personally, do not take this line when considering the income space: making the richest person better off yields more inequality, which might make everyone else a little less happy, suggesting that it is not a pareto improvement in utility space.)

Potential pareto improvements are often defended on the grounds that they might *actually* yield a pareto improvement because the government might *actually* do the transfers. However, that means we are really talking about a pareto improvement, and should present it as such.

The space of cost-benefit analysis is a 'money-equivalent' space—the space of monetized benefit, or willingness-to-pay. If the total willingness-to-pay for a project is positive, then we say that it passes the compensation test, and is 'good' by this criterion.

Willingness-to-pay is also known as 'compensating variation', and is equal to the amount that an individual would pay to secure the project (evaluated as if the project goes forward).

One may also think of willingness-to-pay in terms of opportunity costs. In particular, defining the opportunity cost of a good as its value in its next-best use, willingness-to-pay and opportunity cost are the same thing. For example, what is the CV given up by a person whose labour must be used to create a public good? It is just the opportunity cost of their labour (from their own point of view). Typically, we pay people for their labour, so if we pay the labourer for their labour, then they don't have to give up as much CV. Further, if labour is supplied on a competitive market, we know that the wage will be exactly equal to the opportunity cost, and the worker will be willing to pay exactly the wage to get back one hour of labour time.

The total opportunity cost (across all people) is defined as the social cost of the good. For goods with no externalities in consumption or production, the opportunity cost (willingness-to-pay) is zero for everyone except the transactors, so the social cost is just the individual opportunity cost.

So, compensating variation<===>willingness-to-pay<===>opportunity cost<===>social cost
where <===> gives a link

The Potential Pareto Criterion (PPC) tells you to find *all* the willingness to pay in the economy and add it up. This could potentially be difficult.

For example, if we have a project that uses (as in the textbook) 1000 bricks to that cost $1 each and $500 worth of labour to build a public toilet (a brick one I suppose) and which produces directly $2000 worth of willingness-to-pay for its use (which is free), then you would proceed as follows:

Direct costs: $1500 for the bricks and labour

This must be raised as tax revenue, but since the PPC that a dollar to any person is valued like a dollar to any other, it does not matter *who* pays the taxes, only the total matters.

Direct benefits: $2000 worth of willingness-to-pay

Indirect costs: somebody had to give up the bricks. what was their (negative) willingness-to-pay to give up the bricks minus the amount that they got paid for the bricks?

If the producer is a monopolist, then the producer sacrificed less than $1000 to produce the bricks, so they were paid more than their (negative) willingness-to-pay to give up the bricks. Eg, the bricks might have had a total cost in terms of all inputs of $500 for the monopolist, yielding a negative $500 indirect cost (or, alternatively, positive $500 benefit).

Alternatively, if the input producers are competitive, and the project is small in the sense that the purchase does not affect the competitive price, then the marginal cost to the producer equals the price, which means that the producers' negative willingness-to-pay is exactly balanced by what they were paid.

somebody had to give up the labour hours. what was their negative willingness-to-pay to give up the labour hours minus the amount that they got paid?

If the labourer(s) are in a union, then this would be like a producer with market power—their negative willingness-to-pay would be smaller than what they got paid, so they would be sitting on a piece of benefit.

If the labourer(s) are in an immobile position with the employer (as with a monopsonistic employer), this would yield the opposite situation.

If the labourer(s) are in a competitive labour market, and the amount of labour purchased for the project is small enough to leave the competitive price unaffected, then the marginal cost to the labourer exactly equals the price. This implies that the labourer's (negative) willingness-to-pay for the labour hours is exactly balanced by what they were paid, so there is no indirect cost or benefit to consider.

Once you hit a competitive market, the buck stops on indirect costs.

Cost Functions and Willingness to Pay

We say that the objective function in cost-benefit analysis is the sum of money and monetized utility across all people. Maximizing this objective function is essentially equivalent to using the potential pareto criterion to construct social indifference curves which rank social alternatives.

This is pretty easy for social alternatives which differ only in incomes—choose the one with the most income.

It is harder if the social alternatives differ in other ways. Consider social alternatives which can be described individualistically by the income of each person, *y*, the prices faced by each person, *p*, and a set of characteristics of the social alternative which are relevant to the person, *z*. Here, *y*, is a single number, *p* is a set (vector) of all prices for all commodities, and *z* is a set (vector) of information about everything else that matters to the person. Thus, *z* could be quite complex indeed. It would include some things about the person, like the type of household they live in (how many people, how old), the type of environment they live in (how polluted, how fun) and so on. Most of these characteristics would be the same in the different social alternatives being considered, but some will differ.

You recall the utility function, which gives the utility level associated with each level of income—*u(y)*.

Its inverse is called the cost function—$\tilde{c}(u) = u^{-1}(y)$, and it gives the amount of money necessary to achieve any given level of utility. There is a tilda on *c* because I want it to depend on more arguments below.

The utility function must be different at every price situation because if things cost more, then we are less happy with the same amount of money. Thus, we should really write *u* as a function of at least *p,y*. Further, *u* should depend on other things in a social allocation—if you live in a big household, money doesn't go as far as if you live alone. If you are old, you might need more money to achieve a given level of income, likewise if you are disabled. So, *u* should depend on *z*, too. So, we write *u=V(p,y,z)* and call this thing the "indirect utility function". It is indirect because it tells the amount of utility you get when you maximize utility subject to the constraint of having income *y* to spend on goods which cost *p* in an environment defined by *z*.

We can invert this thing around *y* to get a cost function. Define the cost function as $C(p,u,z) \equiv V^{-1}(p,y,z)$ where the inverse is around *y* and where

$V(p,y,z) = \max_q u(q)\, st \sum_{j=1}^{M} p^j q^j = y$, is the maximum utility attained when choosing commodities q subject to a budget constraint.

How Do We Use the Cost Function to Monetize Non-Income Differences Between States?

If we knew the cost function, *C* (or its inverse, *V*), we could monetize the difference between social states as follows.

Define *consumer surplus* as a measure of the monetized difference in utility between social states, and denote it *s*. You could get a money measure of this difference in lots of different ways. We will consider one way in detail, the method of willingness to pay.

Define *willingness-to-pay* (sometimes called *compensating variation*) as the amount you'd be willing to pay to secure the project. One can fully describe the social allocations *0* and *1* with the incomes, prices and other characteristics of those states, denoted $y_0, y_1, p_0, p_1, z_0, z_1$. Index these pieces of information and the surplus *s* with superscript *i* for each person.

If only incomes varied between the social states, obviously, you'd be willing to pay $s^i = y_1^i - y_0^i$.

If more than just income varied, you'd have to solve the problem: how much could I pay out of income in state 1 such that I'd be just as happy as I am in state 0? This is given by solving $V(p_0^i, y_0^i, z_0^i) = V(p_1^i, y_1^i - s, z_1^i)$ for *s*.

Since we know the cost function is the inverse of *V*, we can write the solution to this as $C\left(p_1^i, u_0^i, z_1^i\right) = y_1^i - s^i$ where $u_0^i \equiv V(p_0^i, y_0^i, z_0^i)$ is the utility level in state 0.

$s^i = y_1^i - C\left(p_1^i, u_0^i, z_1^i\right)$—willingness-to-pay is the amount by which your income in the new state exceeds the amount you'd need in the new state to get the utility level of the old state.

Cost benefit analysis says add up the willingness-to-pay across all people and if it exceeds zero, the project is good—the maximand of cost-benefit analysis is therefore

$$\sum_{i=1}^{N} s^i = \sum_{i=1}^{N} y_1^i - \sum_{i=1}^{N} C(p_1^i, u_0^i, z_1^i)$$

(If prices and characteristics are unchanged, the left-hand piece of this is equal to y_0^i for all people. In this case, the maximand reduces to the sum of the income changes.)

This framework is very general. Once you know the cost function, you have everything you need to do cost benefit analysis for any project which affects incomes, prices, pollution, household structures, anything. The trick then is to find out what the cost function is.

How Do We Learn About The Cost Function?

We can use the cost function to monetize the differences in utility between social allocations for each person. The cost function has some pieces that are in principle observable, and some pieces that are in principle not observable (without further assumptions).

Prices

The pieces that are observable are the pieces of the cost function which depend on prices. Consider a single individual and draw an indifference map over quantities of two commodities, q^1, q^2. Now draw a budget constraint on, with slope equal to the price ratio p^2 / p^1 and with intercept on the q^2 axis given by y / p^2. This intercept is the amount of commodity 2 you get if you spend all your money on it. There is also an intercept on the q^1 axis given by y / p^1.

The indirect utility function gives the value label for the highest indifference curve that you can attain in this situation—the indifference curve that is just tangent to the budget constraint.

The cost function is a function of this value label. If you know which indifference curve you are on, and you know the price vector, then you can identify the income it takes to attain that indifference curve. This income is what the cost function gives you, and it is observable on the indifference map.

Cost is equal to the price multiplied by the intercept of the budget line which is just tangent to the indifference curve under consideration.

Consider two social allocations which differ in the prices of goods and in the incomes of people. Willingness-to-pay *s* is given by the difference in cost between the situations *evaluated at the new price vector.*

This version of willingness-to-pay (which uses cost functions without ever defining them) is discussed in Eaton, Eaton and Allen.

Other Characteristics

We can identify the cost of price and income changes because people make choices that reveal their indifference curves. However, we cannot use the same strategy to identify the cost of other characteristics, such as pollution, arctic terns, faraway forests, and loads of other things which are not under our control. In these situations, we have two choices: (1) We can ask people about their willingness-to-pay for these changes; or (2) We can try uncover tricksy situations in which people *do* make choices that reveal the relevant willingness-to-pay.
Read Portnoy, Paul, 1994; Hahnman, W.N., 1994; and Diamond, Peter and Jerry Hausman, 1994 all on my website for further ideas here.

p 99 For a firm, the gross market price typically equals the opportunity cost *for the firm* (it measures what the firm sacrifices).

P 99 For a government, the market price of an input will not give the marginal social cost (that is, opportunity cost) of the input if:

10. there are externalities in production,
    a. for example, positive externalities like technological spillovers, which cause the price to exceed the marginal social cost,
    b. or, for example, negative externalities like pollution in the production of steel, which causes the marginal social cost to exceed the price;
11. the market isn't competitive,
    a. for example, if the seller is monopolistic (as with a perfect union), so that the price exceeds the marginal social cost,
    b. or, for example, if the buyer is monopsonistic (as with a mining firm or a government monopoly occupation like surgeons), so that the marginal social cost exceeds the price;
12. information problems distort the input market,
    a. for example, if there is hidden information (adverse selection) as in the market for used goods, where the market for certain goods may not exist,
    b. or, for example, if there is hidden action (moral hazard) as in the market for work where effort matters, where the market price (which includes an efficiency wage premium) for the good exceeds the marginal social cost (the marginal value of leisure).

p 99 It isn't really stated, but the outputs of government projects are typically not competitively priced, so we need to assess the marginal social *benefit* as well.

p 99 To value things where the market price is misleading, we need *shadow prices*.

p 99 The input constraint example from Chapter 4 showed that the appropriate shadow price is the price of the input that would induce the constrained level of use.

p 102 What is the shadow price of labour? If the labour market is competitive and so forth, then the shadow price is just the market price.

1. What if there is involuntary unemployment,
    a. for example, because firms pay wage premia to lower turnover, or to induce effort from workers, or get better workers, or to get workers to internalise firm goals (gift exchange),
    b. or, because workers are able to raise their own wage through collective (unions) or individualistic rent or quasirent extraction,
    c. or, because there are institutional or customary constraints like minimum wages that limit the scope of market wages.
2. Take for example, a minimum wage environment where there is excess supply of labour, and where people volunteer labour in a zero/one fashion depending on their reservation wage (valuation of leisure/nonmarket activity). The market wage is w3, and n1 workers get the job, though n3 workers want it. Every employed and involuntarily unemployed person has a reservation wage less w3.
3. What is the shadow price of labour here? If workers were allocated to jobs in order of reservation wages (there is no reason to think they would be), then the marginal social cost of labour is w1.

4. Thus, there are bounds on the opportunity cost of labour, probably w1<MSC<w3.

p 105 What do we do about taxes? Taxes in input markets make it so that prices do not reveal the marginal social cost of inputs.

1. You can rewrite eq (8.1) as *gross price - new tax revenue*. This also works for the material on muliple taxation in the appendices.
2. The intuition that works for me is as follows. The gross price is the marginal social value (opportunity cost) for new uses of the same input. However, if some new input is produced on behalf of the project, then new tax revenue is generated. This new revenue has no welfare effect (because distortions and deadweight loss are ignored so far), so it should be subtracted off the calculation of total social costs.
    a. Thus, if all the input is coming from new production, then all the input generates new tax revenue, all of which should be subtracted off. This results in using the supply price (which is net of taxes) times the quantity as the social cost.
    b. On the other hand, if all the input comes from old production, there is no new tax revenue, and nothing is subtracted off. This results in using the consumption price (which includes taxes) times the quantity.
    c. So, what you need to do is *predict* the change in tax revenue that the project causes, and subtract this off of the gross cost of all the inputs.
    d. Another rationalisation is this: if taxes are *not very* distortionary and input markets are competitive and there are no externalities, then opportunity cost of inputs really is almost equal to the gross cost minus net new tax revenue.

p 107 Okay, so taxes are easy. What about if the supply price is not the opportunity cost of the producer? For example, monopolists charge more and produce less. Then you need to resort to either accounting or statistical type methods to recover the cost function of the firm.

p 107 It is unstated, that even if the supply price is equal to the opportunity of the producer, this does not mean that the opportunity cost of the producer is equal to the social cost. If there are any externalities this equality doesn't hold.

Chapter 9: Big Projects and Price Changes

p 113 Projects considered up to now didn't change prices.
p 113 define *consumer surplus* (CS) as the measure of how much better or worse you are in one environment than another.
p 113 define *compensating variation* (CV) as the amount the individual would pay to secure the project (which comes with a price and income).
p 114 Consider figure 9.1. Price goes up from p1 to p2, and quantity goes down from q2 to q1 *for an individual*.

1. At worst, this person could need (p2-p1)q1 to be as well of as before. They'd need exactly this much if they were completely unable to substitute to other goods.
2. At best, this person could need (p2-p1)q2 to be as well of as before.
3. So, maybe these provide bounds to the monetized welfare effect of the price change.
4. We could break the price change into a sequence of very small changes and get that the total CV is equal to the funny shape that crawls up the demand curve.

p 118 Define *producer surplus* (PS) as the measure of how much suppliers of factors (such as labour) are better off in one environment than another.

Why should we care about consumer or producer surplus that does not accrue to a person? For example, if the project is 'raise the price of gas', we can estimate the loss in consumer surplus off of the marshallian demand curve, but why should we care about producer surplus—after all, firms aren't people.

Firms are owned by people, and economic profits will be distributed to owners. Thus, producer surplus will end up in the hands of people, so we should count it. The distributional neutrality of 'sum of Cvs' allows us to ignore the incidence of that producer surplus.

This is not true, of course, if the owners of firms are not local.

Aren't firms supposed to be competitive? How can there be an increase in economic profits in the market if there is free entry? If there is free entry and there are no increasing returns to scale, then the long run supply curve must be flat. If the long run supply curve is flat, then there is no change in producer surplus if the government raises the price of gas.

Recall that standard cost benefit is not concerned with the distribution of surplus. Thus, the fact that capital owners are concentrated, so that economic profits accrue to the few is not important. That is why you can cancel out losses to CS with gains to PS.

p 118 It is quite analagous to CS, except that it is drawn from information in the supply curve.
p 119 Figure 9.4 shows a PS measure. PS is typically *positive* when the price goes up.
p 120 Consider the same labour market environment as above, where workers get jobs in order of their reservation wage. If the wage goes up, new people get jobs, paying close to their reservation wages. But the old people just get paid more for what they were willing to do for less. This is the origin of the producers' surplus.
p 127 Compensated (or Hicksian) demand curves capture income effects.
p 128 Figure 9.9 shows how the compensated demand curve traces out the demand as price changes but utility stays constant (rather than income staying constant). The authors draw the compensated demand curve as a straight line between the demand curve at ya and the

demand curve at yb. However, since the compensated demand curve is needed to estimate the ‘income effect inclusive’ value of yb (ya+CV), you cannot draw the marshallian demand curve D(yb) without first knowing the compensated demand curve Dc.

p 129 The authors say that: (1) income effects are small so they don’t matter; and (2) that they are not easy to estimate. Claim (1) may or may not be true, depending on the problem. Income effects are large if the CVs are large or if the demand curve is highly responsive to income. Claim (2) is not true. We can estimate the compensated demand curve via structural modelling of the cost function.

General Notes about Consumer and Producer Surplus

For an individual, consumer surplus is any dollarised measure of that person’s utility difference between the status quo and the project. “Compensating Variation” (aka “Willingness-to-Pay”) is the maximum amount that the individual would pay to secure the project against the alternative of the status quo. “Equivalent Variation” (aka “willingness-to-accept”) is the minimum amount the individual would accept to take the status quo against the alternative of the project.

For an individual, the compensated demand curve is the demand curve which holds utility (rather than income) constant as the price varies. For a project which changes prices, the compensating variation for the price change is given by the area between the status quo price and the project price to the left of the compensated demand curve.

We don’t usually have individual compensated demands; rather we usually have aggregate uncompensated (marshallian) demands. Individual uncompensated demands hold individual income (rather than utility) constant. Aggregate uncompensated demands give the sum of demands across individuals holding the sum of income (total income) constant. In order to believe that this is okay, you have to believe that (1) it is okay to sum up demands; and (2) it is okay to use uncompensated demands; and (3) it is okay to hold total income (rather than the set of individual incomes) constant in the aggregate demand.

It is in fact okay to sum up demands. The reason is that you are interested in the sum of compensating variations, each of which is the area to the left of a demand curve. The sum of the areas is equal to the area of the sum (it is just math).

It is not okay to use uncompensated demands, because it ignores the income effect. However, if you assume that the income effects are “small”, you go ahead. It is not okay to hold total income constant, because it ignores the effect of the distribution of income on aggregate demand. Distribution does not affect demand if and only if rich and poor people spend marginal dollars the same way. However, if you assume that they spend marginal dollars “not too differently”, you go ahead.

OPTIONAL Diversion: Consumer Demand and Consumer Surplus Measurement

define $u / V(p,y)$ as the utility you get from income y when you face prices p.
define $y / C(p,u)$ as the cost (income) of attaining utility level u when you face prices p.
C is the inverse of V around y. That is, $C(p,u) / V^{-1}(p,y)$.

Consumer Surplus and CV

Compensating variation is the amount you'd be willing to pay to secure the project:

$V(p_1, y_1 - CV) = V(p_0, y_0)$ where the subscripts indicate states, and state 0 is the status quo.

Invert around income at p_1 to evaluate CV:

$y_1 - CV = C\left(p_1, V(p_0, y_0)\right)$, so that income less CV equals the cost at p_1 to get the utility of state 0.

You can rewrite it as

$$CV = y_1 - C\left(p_1, V(p_0, y_0)\right)$$

or

$$CV = y_1 - y_0 - \left(C\left(p_1, V(p_0, y_0)\right) - C\left(p_0, V(p_0, y_0)\right)\right)$$

Here, CV is the increase in income less the cost of the price change.

So, if we knew the cost/utility function, we would know CV for any project which changes prices and income. It turns out we can estimate the cost/utility function up to a monotonic transformation of u. That is enough, because we only need to know cost at one level of utility ($u_0 / V(p_0,y_0)$) to estimate CV.

Estimating Cost

Shepphard's Lemma:
$$\frac{\partial C(p,u)}{\partial p^j} = h^j(p,u)$$

where $j=1,...,M$ indexes the M commodities, and $h^j(p,u)$ is the quantity demanded for the *j'th* good given prices p and utility u. We call $?^j$ the *compensated demand* for good j.

Logarithmic Shepphard's Lemma:
$$\frac{\partial \ln C(p,u)}{\partial \ln p^j} = w^j(p,u)$$

where $j=1,...,M$ indexes the M commodities, and $w^j(p,u)$ is the share of expenditure commanded by the *j'th* good given prices p and utility u. We call $?^j$ the *compensated expenditure share* for good j.

These two versions of Shepphard's Lemma tell us how cost is related to demand. Imagine that you went to 4 movies per month and the price of a movie went up by a dollar. How much would your costs rise? Your intuition might tell you that costs rise by $4 per month. The impact of a dollar price change is just the quantity you demand. Shepphard's Lemma tells you that this is generally true.

Consider the logarithmic version of Shepphard's Lemma. Imagine that you spend half your income on rent and the price of rent goes up by 10%. What does your intuition tell you about how much money you'll need to be as well off? Hopefully, it tells you that you need ½ * 10%=5% more money (ie. the elasticity of cost is the expenditure share). Shepphard's Lemma states that this intuition is true for all price changes.

Sadly, u is not observable. Since V and C are inverses of each other, if we know C, we know V. So, we can substitute V into $w^j(p,u)$ to create the marshallian (or uncompensated) expenditure share $w^j(p,y) = w^j(p, V(p,y))$, which depends on observables p and y. Alternatively, substitute V into $h^j(p,u)$ to create the marshallian (or uncompensated) demand $h^j(p,y) \equiv h^j(p, V(p,y))$, which depends only on observables p and y.

We know that CV is the difference between cost functions. Shepphard's Lemma tells you that demands are equal to the derivative of cost functions. Thus, we could *integrate* compensated demand functions to get the difference between cost functions (CV). In particular,

$$CV = C(p_1,u) - C(p_0,u) = \int_{p_0^j}^{p_1^j} \frac{\partial C(p,u)}{\partial p^j} = \int_{p_0^j}^{p_1^j} h^j(p,u)$$

CV for a change in the price of a single good j is equal to the integral of the j'th compensated demand equation.

This is the link to those strips on demand equations. However, the strips in demand equations are strips in aggregate marshallian demand equations. The integral of aggregate compensated demand coincides with the sum of integrals of micro compensated demands if and only if preferences are quasihomothetic (see below). Compensated demand *never* coincides with marshallian demand because compensations must be 'spent' somewhere (but people—such as Sugden and Williams—claim that the deviations are 'small').

Two Examples of Parametric Cost Functions and Demand Systems

Gorman Polar Form: Quasihomothetic Preferences

Preferences are q*uasihomothetic* if cost can be written as $C(p,u)=A(p)+B(p)f(u)$ for some A and B. Consider

$$C(p,u) = \prod_{k=1}^{M} \left(p^k\right)^{a^k} + \sum_{k=1}^{M} b^k p^k f(u)$$

where $\sum_{k=1}^{M} a^k = 1, \sum_{k=1}^{M} b^k = 1$, and f is any monotonically increasing function of u.

Inverting around u, we get the (indirect) utility function:

$$f\left(V(p,y)\right)=\frac{y-\prod_{k=1}^{M}\left(p^k\right)^{a^k}}{\sum_{k=1}^{M}b^k p^k}.$$

Taking the derivative of cost, we get the following compensated demands:

$$h^j(p,u)=\frac{a^j}{p^j}\prod_{k=1}^{M}\left(p^k\right)^{a^k}+b^k f(u)$$

We know (from the fact we derived it) that the integral of the compensated demand is the cost function. So, if you draw the compensated demand equation, you find that given utility, it is almost linear in the reciprocal of own-price. Integrating it is just measuring the area covered as the price changes.

Substituting for *f(u)=f(V(p,y))*, you get marshallian demands:

$$h^j(p,y)=\frac{a^j}{p^j}\prod_{k=1}^{M}\left(p^k\right)^{a^k}+b^k\left(\frac{y-\prod_{k=1}^{M}\left(p^k\right)^{a^k}}{\sum_{k=1}^{M}b^k p^k}\right).$$

or

$$h^j(p,y)=a^j\bullet\frac{1}{p^j}\prod_{k=1}^{M}\left(p^k\right)^{a^k}+b^k\bullet\left(\frac{y}{\sum_{k=1}^{M}b^k p^k}-\frac{\prod_{k=1}^{M}\left(p^k\right)^{a^k}}{\sum_{k=1}^{M}b^k p^k}\right)$$

The marshallian demands are 'almost' linear with just two terms: the reciprocal of own-price, and income. However, it is clear that marshallian demands do not coincide with compensated demands because, although the first part is the same, the second part is not the same, and the second part contains prices.

If you have data on quantities, prices and incomes for a bunch of individuals, you can estimate the parameters a^k and b^k. If you estimate the parameters, you have everything you need to create the function *C*. That means that you have everything you need to estimate CV's for everyone in your population.

These CV's are exact. You don't need to assume that income effects are 'small', or that marshallian and hicksian demands coincide, or that aggregate demand reveals the same information as micro-demand.

These demand equations are linear in income given prices. This linearity implies that for any aggregate level of income, demand is the same, because marginal dollars are spent identically by everyone. This implies that aggregate demand *does* in fact reveal the same information as micro-demand. This is the first nice thing about quasihomothetic preferences.

Quasihomothetic preferences are very special for the purposes of cost-benefit analysis. In particular, if demands are quasihomothetic, then

(1) cost can be written as $C(p,u)=A(p)+B(p)f(u)$ where A and B depend on prices (Gorman);
(2) aggregate demand does reveal the same information as micro-demand (write out the aggregate demand, you'll see);
(3) the distribution of income does not affect aggregate demand (so the Boadway paradox doesn't bite);
(4) the sum of CV gives a full ordering over alternatives which is independent of which we take as the 'base' alternative (ie. CV ranking=EV ranking)

The unfortunate thing about quasihomothetic preferences is that it looks like people don't have them. Rich people spend their marginal dollars differently from poor people.

Almost Ideal demand system (Deaton and Muelbauer)

Consider

$$\ln C(p,u) = \sum_{k=1}^{M} a^k \ln p^k + \frac{1}{2}\sum_{k=1}^{M} \sum_{l=1}^{M} a^{kl} \ln p^k \ln p^l + \prod_{k=1}^{M} \left(p^k\right)^{b^k} f(u) \tag{1}$$

where $\sum_{k=1}^{M} a^k = 1, \sum_{k=1}^{M} b^k = 0, \sum_{k=1}^{M} a^{kl} = 0 \ \forall l$, and f is any monotonically increasing function of u. Inverting around u, we get the (indirect) utility function:

$$f\left(V(p,y)\right) = \frac{\left(\ln y - \sum_{k=1}^{M} a^k \ln p^k - \frac{1}{2}\sum_{k=1}^{M} \sum_{l=1}^{M} a^{kl} \ln p^k \ln p^l\right)}{\prod_{k=1}^{M} \left(p^k\right)^{b^k} f(u)} \tag{2}$$

Taking the elasticity of cost, we get the following compensated expenditure shares:

$$w^j(p,y) = a^j + \sum_{k=1}^{M} a^{jk} \ln p^k + b^j \prod_{k=1}^{M} \left(p^j\right)^{b^k} f(u)$$

and substituting for $f(u)=f(V(p,y))$, you get expenditure shares

$$w^j(p,y) = a^j + \sum_{k=1}^{M} a^{jk} \ln p^k + b^j \left(\ln y - \sum_{k=1}^{M} a^k \ln p^k - \frac{1}{2}\sum_{k=1}^{M} \sum_{l=1}^{M} a^{kl} \ln p^k \ln p^l\right) \tag{3}$$

Here, expenditure shares are almost linear in log-prices and log-income. Here, the expenditure shares rise or fall by the same amount for a given proportionate increase in income, whether or not the person is rich or poor. This is restrictive, but probably not as wrong as assuming quasi-homotheticity.

Table: Almost Ideal demand system estimates:
Canada, 5 regions, 12 years, 60 price vectors, 19276 households
No Demographic Effects

Log of Likelihood Function = 296078.
Number of Observations = 19276

Parameter		Estimate	SE	Parameter	Estimate	SE
A1	Food In	0.206	0.001	A27	-0.007	0.003
A2	Food Out	0.072	0.001	A28	0.002	0.003
A3	Rent	0.376	0.001	A29	-0.019	0.003
A4	HH Oper	0.079	0.001	A33	0.104	0.005
A5	HH Furn	0.034	0.001	A34	0.004	0.002
A6	Ad Cloth	0.077	0.001	A35	-0.043	0.002
A7	Ch Cloth	0.007	0.000	A36	-0.035	0.002
A8	Car Oper	0.080	0.001	A37	-0.002	0.001
A9	Pub Trans	0.036	0.001	A38	-0.007	0.003
B1	Food In	-0.045	0.001	A39	0.042	0.002
B2	Food Out	0.027	0.001	A44	0.002	0.008
B3	Rent	-0.135	0.002	A45	-0.035	0.007
B4	HH Oper	0.006	0.001	A46	-0.006	0.005
B5	HH Furn	0.034	0.001	A47	-0.002	0.003
B6	Ad Cloth	0.048	0.001	A48	-0.026	0.003
B7	Ch Cloth	0.006	0.000	A49	0.025	0.003
B8	Car Oper	0.056	0.001	A55	-0.049	0.009
B9	Pub Trans	-0.003	0.001	A56	0.012	0.005
A11		-0.138	0.012	A57	0.035	0.004
A12		0.048	0.008	A58	-0.016	0.003
A13		-0.100	0.004	A59	-0.025	0.002
A14		0.029	0.008	A66	0.040	0.005
A15		0.088	0.007	A67	0.007	0.003
A16		0.057	0.006	A68	-0.005	0.003
A17		-0.003	0.003	A69	0.006	0.002
A18		0.032	0.005	A77	-0.021	0.004
A19		-0.033	0.004	A78	-0.002	0.001
A22		-0.039	0.007	A79	-0.001	0.001
A23		0.048	0.003	A88	0.017	0.004
A24		0.018	0.005	A89	0.008	0.002
A25		0.005	0.006	A99	-0.002	0.002
A26		-0.053	0.004			

Equation (3) may be estimated by adding an error term to the right hand side and choosing parameters to minimise the sum of squared errors. All you need is micro-data on household spending on various commodities, the prices of those commodities and the total expenditure of each household. The table above gives parameter estimates for Canada 1969-1999 for the almost ideal demand system with no demographic effects. Prices are normalised so that the price of every good in Ontario 1986 is equal to the total expenditure of the average single childless adult. This implies that rightmost term in the expenditure share equation is zero for the average single childless adult in Ontario in 1986, and thus makes it easier to interpret the parameters.

Now consider including restaurant food in the PST base, so that the price of 'food out' goes up by 7.5%. What is the CV for each household? Every household has cost depending on *f(u)* given by (1). Every household has an *f(u)* value given by (2). The CV calculation holds utility constant, and thus holds *f(u)* constant.

If the b^j parameters were all zero, then *lnC* would change by the same amount for everyone:

$$a^2 \bullet 0.075 + \frac{1}{2}\sum_{k=1}^{M} a^{2k} \bullet 0.075 \bullet \ln p^k = 0.072 \bullet 0.075 + \frac{1}{2}\sum_{k=1}^{M} a^{2k} \bullet 0.075 \bullet \ln p^k \approx 0.005 .$$

Since the b^j parameters are not zero, the proportional impact on cost is different for rich and poor, depending on whether or not the last term in the cost function rises or falls with the price of food. The estimated $b^2 = 0.027$, is positive so that the difference in *lnC* must rise with utility *f(u)*. Thus, the estimated proportional change in expenditure necessary to hold utility constant over this increase is larger for the rich than the poor.

END OF OPTIONAL DIVERSION

Chapter 10: Indirect Effects of Price Changes (General Equilibrium Considerations)
Also read introductory section to Blackorby and Donaldson Mathematical Social Science paper, on my website.

p 134 We have been considering single markets up to now. In general, it is not possible to change a price in one market without changing prices in other markets due to the fact that consumers and producers can substitute in the face of the price change.

p 134 General Equilibrium (indirect) effects are hard to predict. You need a general equilibrium model. Only the simplest GE models are analytically tractable. In more complex cases, you'll need a computable GE simulation.

p 134 Even with good prediction, aggregation of all effects can be difficult. Which effects (CV changes) should be counted, and which should be considered as part of the direct effect?

p 135 An example of effects that should not be counted is in Figure 10.1a and 10.1b. Closing a rail service (which was provided at zero cost to the producer) must neccessarily hurt consumers. It reduces their opportunity sets. In Figure 10.1a, you see this loss as area ABC, which covers the surplus lost as we raise the price to beyond what anyone is willing to pay.

p 135 People then substitute. They increase demand for bus service (which is costless to the provider), pushing demand out to Db(pr') in Figure 10.1b. It looks like there is increased CV here, from area DEF to area DGH, and it looks like the increase in Figure 10.1b is larger than the loss in Figure 10.1a. However, intuitively, we know that this cannot be so.

p 136 The change is surplus comes only from the direct effect.

p 136 You might ask, what if there was an extremely good substitute for the closed rail service. That is, what if the bus service was just as cheap and almost as fast as the rail service. Then, people would substitute to the bus service, and not lose very much. This would be reflected in the demand curve for the rail service as highly elastic (flat) demand, because if the rail company raised its price even a little, it would lose all its customers. The fact that demand is not so elastic in Figure 10.1a means that the bus service substitute is not a very good one for lots of people.

p 137 What if the bus service supply curve wasn't flat? What if the increased demand resulted in a price increase in the bus market? Or what if it resulted in a price change in some other market, say the housing market?

p 139 The market for rail trips and tenancies is in Figures 10.2a and 10.2b. **Instead please use the figure from Blackorby and Donaldson 1999 on my webpage.** The discussion of pages 140-142 are drawn from Blackorby and Donaldson rather than Sugden and Williams. Both give the right result, but Sugden and Williams' argument is wrong. The Blackorby Donaldson figure is an example where the price of rail travel goes *down*.

p 140' There is a price *reduction* in the price of rail trips, which pushes demand *up* initially along $D_r(p_r, p_h^b)$, where p_r refers to the price of rail trips and p_h^b refers to the price of housing before the rail price change. At first blush, we would integrate this demand curve from point C to point F and be done.
However, we know that with more rail traffic, there will be greater demand for suburban housing, and the price for housing will rise to p_h^a (price of housing after). But since

tenancies and train trips are complements, this price rise pushes *down* the demand for train trips to $D_r(p_r, p_h^a)$. Thus, the quantity demanded in the rail market will go from r^a to r^b that is from point C on $D_r(p_r, p_h^b)$ to point E on $D_r(p_r, p_h^a)$ in this general equilibrium analysis.

p 140' What is consumer surplus in the rail market? There are two ways to start. Either do the rail market first, or the housing market first.

(1) If we do the rail market first, then we calculate gain to CS off of $D_r(p_r, p_h^b)$ (because the housing price change hasn't passed through yet), and get ACFD. Then, we go to the housing market. The rail price has already dropped, so we look at $D_h(p_r^a, p_h)$ to get the increase in the price of housing from p_h^b to p_h^a. We get the loss to CS off of $D_h(p_r^a, p_h)$, which is GILJ,and a gain of PS of GIKJ, which yields a total surplus decrease in the housing market of ILK (CS loss exceeds PS gain). Thus, at first blush, the total is ACFD in the intervention market minus ILK in the secondary competitive market.

(2) Alternatively, we start in the housing market. Here, the price goes from p_h^b to p_h^a, and we use the demand curve $D_h(p_r^b, p_h))$ which uses the initial rail market equilibrium price. CS goes down by GHKJ, but PS goes up by GIKJ, which yields excess PS in the housing market of (positive) HIK. Then, we go to the rail market, and compute CS off of the demand curve which uses the new housing price pha, $D_r(p_r, p_h^a)$, yielding CS of ABED. Thus, at first blush, the total change in surplus is ABED plus HIK, which could be different from what we got above.

(3) Another approach is to use these two methods to sandwich the effect of a small part of the price change in the rail market. In the intervention market, we will get areas sandwich'ed between integrals of the before and after demand curves. In the secondary competitive market, we will get triangles like HIK or ILK. As the price changes get small, the triangles will go to zero, but the integrals in the intervention market will go to nonzero rectangles.

   (1) Consider the first part of the price change. The triangles in the secondary competitive market are near zero, and the rectangle in the primary market is AC. After this small piece of price change, the demand curve D_r moves part of the way from $D_r(p_r, p_h^b)$ to $D_r(p_r, p_h^a)$. Thus, we should add up the CS along a demand curve that goes from E to C, yielding CS of ACED.

   (2) Thus, you need to integrate a *reduced form* demand curve in the intervention market that traces out the demand for rails taking into account price changes in secondary markets. That is, let $p_h = P_h(p_r)$ and substitute it into $d_r, d_r = D_r\left(p_r, P_h(p_r)\right) = \widetilde{D}_r\left(p_r\right)$. This reduced form demand traces

out the equilibrium path. So, you integrate $\widetilde{D}_r(p_r)$ from p_r^b to p_r^a. Since the secondary competitive market results in exactly offsetting producer and consumer surplus along the equilibrium path, you need not consider it.

(3) Sometimes people say that the reason you integrate from C to F is that IKL=CEF. This is not a reason. It is a consequence of the fact that in the general equilibrium, producer and consumer surplus are exactly offsetting along the equilibrium path in all secondary competitive markets.

General Notes about General Equilibrium

There are two kinds of effects to consider: (1) the direct effect of the intervention on the intervention market, and the indirect effect of the intervention on other markets through the price mechanism; and (2) *feedback* effects from all those other markets *back* to the intervention market. These feedback effects are called "general equilibrium effects".

Cost-benefit analysis proposes tricks for dealing with both (1) and (2). The trick for (1) is to stop counting direct and indirect effects when you hit a competitive market. The reason is that once you hit a competitive market, changes in consumer and producer surplus exactly cancel out, so that their sum is exactly zero.

The trick for (2) is if *all* of the other markets (other than the intervention market) are competitive, then the above holds in all the other markets (so, you don't have to worry about the sum of consumer and producer surplus in those markets), *and* the feedback effect onto the intervention market is much simplified. In particular, will include all the feedback effects if you use the area to the left of the demand curve which holds *only* the own price constant and allows all other prices to be determined by the own price. This is sometimes called the 'reduced form' demand.