CHOOSING TREATMENT POLICIES UNDER AMBIGUITY

Charles F. Manski

Northwestern University

Economists studying choice with partial knowledge assume that the decision maker places a subjective distribution on unknown quantities and maximizes expected utility.

Someone lacking a subjective distribution faces a problem of choice under ambiguity.

Consider treatment choice for a population.

Ambiguity arises when a planner has partial knowledge of treatment response and does not feel able to place a subjective distribution on the unknowns.

One may first eliminate dominated policies and then use some *reasonable* criterion to choose an undominated policy.

Sources

Manski C. 2000. Identification Problems and Decisions under Ambiguity: Empirical Analysis of Treatment Response and Normative Analysis of Treatment Choice. *Journal of Econometrics*. 95: 415–442

Manski C. 2006. Search Profiling with Partial Knowledge of Deterrence. *Economic Journal*. 116: F385-F401

Manski C. 2009. Diversified Treatment under Ambiguity. *International Economic Review*. 50: 1013-1041

Manski C. 2010. Vaccination with Partial Knowledge of External Effectiveness. *Proceedings of the National Academy of Sciences*. 107: 3953-3960

Manski, C. 2011. Choosing Treatment Policies under Ambiguity. *Annual Review of Economics*, 3: 25-49.

VACCINATION WITH PARTIAL KNOWLEDGE OF EXTERNAL EFFECTIVENESS

(*PNAS*, 2010)

The problem of choosing an optimal vaccination policy for a population susceptible to infectious disease has drawn considerable attention.

Research studying optimal vaccination has assumed the planner knows how vaccination affects illness rates.

There are two reasons why a planner may have only partial knowledge of the effect of vaccination on illness.

He may not know the *internal* effectiveness of vaccination in generating an immune response that prevents a vaccinated person from become ill or infectious.

He may not know the *external* effectiveness of vaccination in preventing transmission of disease to members of the population who are unvaccinated or unsuccessfully vaccinated.

Knowledge of external effectiveness is most problematic.

A standard randomized clinical trial enables evaluation of internal effectiveness.

However, the outcome data only reveal the external effectiveness of the chosen vaccination rate.

Outcomes with other vaccination rates are counterfactual.

To cope with the absence of empirical evidence, researchers have used epidemiological models to forecast the outcomes that would occur with counterfactual vaccination rates.

Authors typically do not know the accuracy of their assumptions about individual behavior, social interactions, and disease transmission.

Hence, it is prudent to view their forecasts more as computational experiments than as accurate predictions of policy impacts.

I study choice of vaccination policy when a planner has partial knowledge of the external effectiveness of vaccination.

I take the planner's objective to be minimization of the social cost of illness and vaccination.

I consider a simple scenario in which the population is composed of observationally identical persons.

The planner knows that vaccination is fully effective internally. Thus, vaccinated persons never become ill.

Regarding external effectiveness, the planner knows only that the rate of illness among unvaccinated persons decreases as the vaccination rate rises.

The planner observes the illness rate under a policy that vaccinates an observed fraction of the population.

Dominated Vaccination Rates

I show that the empirical evidence and the assumption of monotone external effectiveness imply that certain vaccination rates are strictly dominated.

That is, there exist other vaccination rates that yield lower social cost whatever the true external effectiveness of vaccination may be.

Broadly speaking, low (high) vaccination rates are dominated when the cost of vaccination is low (high)

Minimax and Minimax-Regret Vaccination Rates

How might the planner choose among the undominated vaccination rates?

I derive the minimax and minimax-regret vaccination rates.

These criteria protect against poor outcomes, but in different ways.

The former chooses an action that minimizes maximum cost across all feasible states.

The latter chooses an action that minimizes maximum regret. The regret of a specified vaccination rate in a given state of nature is the cost of this rate minus the cost of the best possible rate.

I refer to minimax and minimax-regret as “reasonable” decision criteria, not as “optimal” ones.

There is no uniquely correct way to choose among undominated actions.

The crux of the problem in choice under ambiguity is that the planner does not know which action is best.

Related Planning Problems

The analysis demonstrates how one may address a class of problems where a planner observes the outcome of a status-quo policy and partially extrapolate to proposed policies.

Manski (2006) studied the criminal-justice problem of choosing a rate of search for evidence of crime, when a planner has partial knowledge of the deterrent effect of search on the rate of crime commission.

I considered a planner who wants to minimize the social cost of crime, search, and punishment. The planner observes the crime rate under a status-quo search rate and assumes that the crime rate falls as the search rate rises.

This planning problem is similar to the vaccination problem.

DIVERSIFIED TREATMENT UNDER AMBIGUITY

(*IER*, 2009)

Suppose that a planner can treat persons differentially.

Examples: medical treatments, sentencing of offenders, active labor-market programs.

He may make a *singleton* allocation, assigning all observationally identical persons to the same treatment.

He could choose a *fractional* allocation, randomly assigning positive fractions of these persons to both treatments.

Portfolio choice in finance has long been framed as a choice among fractional allocations.

Planning has been studied as choice of a singleton.

Fractional allocations cope with ambiguity through diversification.

Diversification enables a decision maker to balance two types of potential error.

A Type A error occur when treatment a is chosen but is actually inferior to b, and a Type B error occurs when b is chosen but is inferior to a.

The singleton allocation assigning everyone to treatment a entirely avoids type B errors but may yield Type A errors, and vice versa for singleton assignment to treatment b.

Fractional allocations make both types of errors but reduce their potential magnitudes.

Suppose that

* there are two feasible treatments.
* treatment is individualistic.
* welfare is a linear function of individual outcomes.
* persons are observationally identical.

This setting eliminates external effects, risk aversion, and profiling as reasons for differential treatment of persons.

FINDING: The minimax-regret criterion always yields a fractional allocation when the better treatment is not known.

Adaptive Diversification

Suppose that a planner must choose treatments for successive cohorts of a population.

Then learning is possible, with observation of the outcomes experienced by earlier cohorts informing treatment choice for later cohorts.

Diversification generates randomized experiments yielding outcome data on both treatments.

I suggest the *adaptive minimax-regret* (*AMR*) criterion, which applies the static minimax-regret criterion using the information available at the time.

Conclusion

Optimal policy choice under ambiguity is not achievable, but reasonable choices based on coherent decision-theoretic principles are achievable.

Planners should not seek to hide ambiguity behind untenable assumptions. They should face up to it.

First study dominance to eliminate clearly bad policies.

There is no objectively correct way to choose an undominated policy. One might maximize subjective expected welfare, maximize minimum welfare, minimize maximum regret, or use another criterion.

Analysis can inform policy choice by characterizing the properties of alternative criteria in specific settings.

IDENTIFICATION OF INCOME-LEISURE PREFERENCES AND EVALUATION OF INCOME TAX POLICY

(working paper, 2012)

Prediction of the response of labor supply to income taxation has long been an important concern of economic policy analysis.

Discussion has been characterized by dueling certitudes.

Conservatives assert that labor supply would increase significantly if tax rates would be reduced. Liberals assert that labor supply is relatively insensitive to tax rates.

The Theory of Labor Supply

Standard economic theory does not predict the response of labor supply to income taxation.

As tax rates increase, a person may rationally decide to work less, work more, or not change his labor supply at all.

A simple model suffices to show this.

Consider a person with a predetermined wage. He allocates a unit of time between paid work and leisure. His net income equals his gross income minus his income tax.

Economists generally suppose that persons prefer to have more income and more leisure.

The essence of the time-allocation supply problem is that a person cannot simultaneously increase income and leisure. Income increases with the amount of time worked and leisure decreases.

Standard theory supposes that the person attaches a utility to each feasible (net income, leisure) pair and chooses time allocation to maximize utility.

Beyond the presumption that net income and leisure are both desirable, theory is silent on preferences.

Different preferences imply different relationships between taxation and labor supply.

Consider a person with no unearned income and ask how labor supply varies with the rate of a proportional tax.

A person with *additive utility* works full time when the tax rate is low and not at all when the tax rate is high.

One with *Leontief utility* works more as the tax rate rises.

Cobb-Douglas utility implies that labor supply does not vary with the tax rate.

Other utility functions imply *backward-bending* labor supply functions. Others yield even more complex relationships between tax rates and labor supply.

Theory does not give a privileged status to particular preferences. Moreover, preferences may be heterogeneous.

Empirical Analysis

The silence of theory has long been appreciated.

Robbins (1930):

> "we are left with the conclusion . . . . that any attempt to predict the effect of a change in the terms on which income is earned must proceed by inductive investigation of elasticities."

Mirrlees (1971):

> "The examples discussed confirm, as one would expect, that the shape of the optimum earned-income tax schedule is rather sensitive to the distribution of skills within the population, and to the income-leisure preferences postulated. Neither is easy to estimate for real economies."

Economists have performed numerous empirical studies of labor supply, using two approaches.

Atheoretical Analysis of Treatment Response

Before-and-after studies compare labor supply before and after a change in tax policy.

Cross-sectional studies compare the labor supply of persons living in different tax jurisdictions.

Revealed Preference Analysis of Labor Supply

Researchers performing revealed preference analysis observe the labor supply decisions made by a study population under a status quo tax policy.

To predict labor supply under new policies, researchers invoke the standard assumption that persons allocate their time to maximize utility.

Hence, a person prefers his chosen (net income, leisure) pair over all other feasible pairs.

Hence, observation of labor supply under a status quo tax policy reveals something about preferences.

Researchers place strong assumptions on preferences to obtain point predictions of response to new tax policies.

They assume that labor supply varies monotonically with net wages.

They assume that the response of labor supply to net wage is homogeneous within broad groups.

Neither assumption has a foundation in theory or evidence.

Credibility aside, these assumptions enable prediction of labor supply under new tax policies.

A large body of research was stimulated by the work of Burtless and Hausman (1978).

The literature has been summarized and critiqued by Pencavel (1986), Killingsworth and Heckman (1986), Blundell and MaCurdy (1999), Meghir and Phillips (2010), Keane (2011), and Saez, Slemrod, and Giertz (2011).

Meghir and Phillips (2010):

> "Our conclusion is that hours of work do not respond particularly strongly to the financial incentives created by tax changes for men, but they are a little more responsive for married women and lone mothers. On the other hand, the decision whether or not to take paid work at all is quite sensitive to taxation and benefits for women and mothers in particular."

Saez, Slemrod, and Giertz (2011):

> "With some exceptions, the profession has settled on a value for this elasticity close to zero for prime-age males, although for married women the responsiveness of labor force participation appears to be significant. Overall, though, the compensated elasticity of labor appears to be fairly small."

Keane (2011):

> "My review suggests that labor supply of men may be more elastic than conventional wisdom suggests."

While authors may differ on the magnitude of labor-supply elasticities, they agree on the sign.

The consensus is that increasing tax rates usually reduces work effort.

Considering the effect of a rise in a proportional tax, Meghir and Phillips (2010) write (p. 207):

> "in most cases this will lead to less work, but when the income effect dominates the substitution effect at high hours of work it may increase effort."

Keane (2011):

> "the use of labor income taxation to raise revenue causes people to work less."

Basic Revealed-Preference Analysis

In Manski (2012), I study the predictions yielded by revealed-preference analysis when one strips away the strong assumptions made in the literature and makes only the two most basic assumptions of standard theory:

(1) a person chooses his time allocation to maximize utility

(2) utility increases with income and leisure (more is better).

I find that combining these assumptions with observation of a person's time allocation under a status quo tax policy does not yield point predictions of labor supply under new policies—one can at most obtain interval predictions.

Moreover, one cannot predict whether labor supply would increase or decrease in response to changes in tax policy.

I find that a precondition for basic revealed preference analysis to have any predictive power is that the status quo and new tax schedules cross at least once.

That is, one policy must yield lower net income than the other at some time allocations but higher net income at other time allocations.

This implies that basic analysis has no predictive power when the status quo is a proportional tax and the new policy is a proportional tax at a different rate.

On the other hand, the analysis may have some predictive power when one policy is a progressive tax and the other is a proportional tax.

Illustration: Labor Supply under Progressive and Proportional Taxes

Let the status quo tax income at rate 15 percent up to $50,000 per year and at rate 25 percent above $50,000.

Consider a new policy that taxes all income at 20 percent.

The two schedules cross when gross annual income equals $100,000, where both take $20,000 tax and yield net income $80,000.

Consider a person with no unearned income whose gross annual income for full-time work is $150,000.

This person earns net income $80,000 under both policies if he works 2/3 of the year and takes 1/3 as leisure.

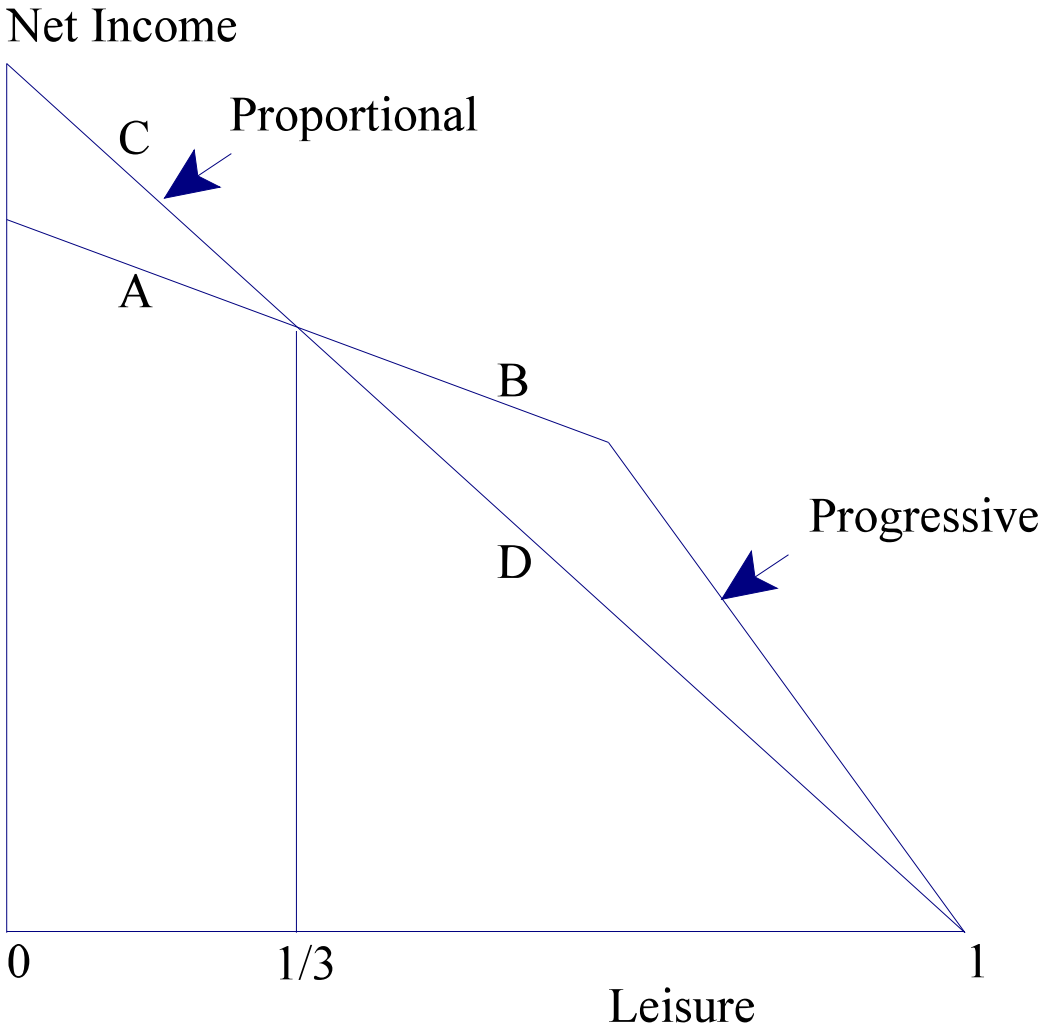
Net Income
Proportional
C
A
B
Progressive
D
0
1/3
1
Leisure

Assumptions on the Preference Distribution

I next explore the identifying power of two assumptions restricting the distribution of income-leisure preferences.

One assumes that groups of persons who face different choice sets have the same distribution of preferences. The other adds restrictions on the shape of this distribution.

The analysis applies methodology developed in Manski (*IER*, 2007).

A Computational Experiment

Given data on labor supply under status quo policy S, the problem is to predict tax revenue per capita under policy T.

I show the predictive power of a sequence of assumptions:

(1) more is better

(2) + persons in specified wage groups have the same distribution of preferences

(3) + preferences have the CES form

(4) + all CES utility functions in a wage group have the same elasticity of substitution.

Let S tax at 20% up to $100,000 and 30% above $100,000. Let T tax at 25%. The schedules cross when gross income equals $200,000.

The population has persons whose wages are ${50, 100, 150, 200, 250, 300, 350, 400} x 1000. The distribution is

P(w = 50) = 0.70, P(w = 100) = 0.20, P(w = 150) = 0.05,
P(w = 200) = 0.02, P(w = 250) = 0.01, P(w = 300) = 0.0075,
P(w = 350) = 0.0075, P(w = 400) = 0.005.

The feasible values of leisure are {0, ¼, ½, 1}.

CES utility functions have the form

$$U_j(Y, L) = [\alpha_j(Y/400{,}000)^{\rho_j} + (1 - \alpha_j)L^{\rho_j}]^{1/\rho_j},$$

$\alpha \in \{0, 0.01, 0.02, \ldots, 0.99, 1\}$

$\rho \in \{-100, -90, \ldots, -20, -10, -1, -0.99, .0, \ldots, 0.99, 1\}.$

Hence, there are $101 \times 211 = 21{,}311$ CES preferences.

The population actually contains persons with 20 distinct CES preferences:

$(\alpha, \rho) \in \{0.25, 0.5, 0.65, 0.75\} \times \{-100, -0.5, 0, 0.5, 1\}.$

The distribution of (α, ρ) is uniform conditional on wage.

Findings

Actual tax revenues under S and T are $7,652 and $9,211.

Knowing only the wage distribution but not labor supply under S, revenue under T lies in the interval [0, $18969].

With data on status quo labor supply, assuming more-is-better yields the bound [$593, $18969].

Combine more-is-better with the assumption that persons with wage up to $200,000 have the same preference distribution and those with wage above $200,000 have the same distribution. Then revenue $\in$ [$3744, $14149].

Next suppose that all persons have CES utility functions. This yields the bound [$6883, $10444].

Finally, assume that the persons in each wage group have the same elasticity of substitution. This yields an empty identification region for the preference distribution within the second wage group. Thus, the assumption of homogeneous elasticity within each wage group is rejected.

Utilitarian Policy Evaluation

A familiar exercise in normative public economics poses a utilitarian social welfare function and ranks tax policies by the welfare they achieve.

This requires knowledge of preferences to
(1) predict tax revenues
(2) compute welfare.

I ask how a utilitarian social planner might compare tax policies when available data and credible assumptions only partially identify population preferences.

Partial knowledge of preferences implies partial knowledge of the welfare function.

The problem is exacerbated if the planner contemplates using tax revenue not only to redistribute income but also to produce public goods.

I approach choice of a tax policy as a problem of planning under ambiguity, as I have done previously when studying other policy choices.

There are multiple difficulties.

1. Basic revealed preference analysis does not predict labor supply under policies that change the production of public goods that directly affect utility.

2. Computation of welfare requires knowledge of the preferences of the population, not just their labor supply.

3. Policies which fix public-good production before observation of labor supply generically yield budget surpluses or deficits, because the ability to predict labor supply is incomplete.

Enriching the Data for Identification of Preferences

We lack the knowledge of preferences necessary to perform credible utilitarian evaluation of income tax policy.

The consensus that increasing tax rates reduces work effort is premature.

Knowledge of preferences for public goods is almost non-existent.

I do not expect that new theory will sharpen our knowledge of preferences.

The only way is to obtain richer data. Some ideas follow:

1. Resume the performance of randomized experiments with tax policy that began with the negative income tax experiments of the 1970s.

Randomized experiments can make it credible to assume that groups of persons who face different tax schedules have the same distribution of income-leisure preferences.

2. Obtain repeated observations from individuals.

Assuming a static model of labor supply, it is useful to repeatedly observe an individual in circumstances with varying wages, unearned income, or tax schedules.

Transitivity implies that basic revealed preference analysis has increasing predictive power as more choices are observed (Samuelson, 1938, 1948; Afriat, 1967).

However, this interpretation of repeated observations rests critically on the static model.

3. Interview persons, pose choice scenarios with various hypothetical wages or tax schedules, and ask them to predict their own choice behavior in these scenarios.