

THE CONTRIBUTIONS OF ROBERT S. PINDYCK

by Ernst R. Berndt, MIT Sloan School of Management

Invited Submission to The Palgrave Companion to MIT Economics

Edited by Robert Cord Springer Nature Switzerland AG

ABSTRACT

Robert S. Pindyck has made significant contributions to a wide and diverse range of economic literatures; summarizing them is challenging. To keep the length of this review manageable, I confine my discussion to four areas in which his contributions have been particularly notable: (i) economic applications of optimal control theory; (ii) market structure and the exploitation of non-renewable resource markets; (iii) uncertainty, irreversible investments, and option values; and (iv) economic issues underlying climate change and climate policy. Professor Pindyck has been a highly acclaimed educator in MIT classrooms, winning numerous teacher awards. He has also been a successful textbook author, with both his microeconomics and econometrics textbooks having been published in multiple editions.

Key Words: optimal control, exhaustible and non-renewable resources, uncertainty, irreversible, option value, sunk costs, climate policy

AUTHOR'S PROFESSIONAL BIOGRAPHY

Ernst R. Berndt is the Louis E. Seley Professor of Applied Economics Emeritus at the Alfred P. Sloan School of Management, Massachusetts Institute of Technology, where he has been a faculty colleague of Professor Robert Pindyck since 1980. Berndt is an elected Fellow of the Econometric Society, served as the first chair of the Federal Economic Statistics Advisory Committee 2000-04, has been a Research Associate at the National Bureau of Economic Research since 1980, and served as Director of its Program on Productivity Measurement. He is the author of the econometrics textbook, ***The Practice of Econometrics: Classic and Contemporary***.

Preliminary and Confidential: Do Not Cite Without Author's Permission

THE CONTRIBUTIONS OF PROFESSOR ROBERT S. PINDYCK

TABLE OF CONTENTS

I.	Professional Biography of Robert S. Pindyck	3
II.	Introduction to Professor Pindyck's Contributions	4
	A. Linear Econometric Models	5
	B. Nomenclature and Notation in Optimal Control Models	6
	C. Pindyck's "Proof of Concept" in Optimal Control Theory	8
	D. Decentralized Optimal Control with Conflicting Objectives	14
III.	Cartels and Natural Resource Exploitation.	18
	A. Game Theory: OPEC as a Sustainable Cartel	18
	B. What Makes Cartels Feasible and Sustainable?	21
	C. Challenging Hotelling on the Pricing of Non-Renewable Resources.	23
	D. Massive Uncertainty in Non-Renewable Resource Markets	25
	E. Measuring Market Power in Dynamic Markets	27
IV.	Uncertainty, Irreversible Investments, and Option Values	29
	A. Irreversibilities Make Net Present Value Criteria Unreliable.	29
	B. Cost vs. Demand Uncertainty and Irreversible Investments	31
V.	Uncertainty Insurance and Irreversibilities in Climate Policy	33
	A. The Critical Importance of Uncertainty in Climate Policy.	33
	B. Climate Insurance to Deal with Massive Uncertainty.	36
	C. Risk Aversion as Another Rationale for Buying Climate Insurance	38
	D. Environmental Protection Expenditure as Irreversible Sunk Costs	39
	E. Insurance Value of Climate Action Dominates Sunk Cost Irreversibilities	40
VI.	Concluding Remarks	40
VII.	Cited References	41

I. PROFESSIONAL BIOGRAPHY OF ROBERT S. PINDYCK

Robert Pindyck arrived on the MIT campus as an undergraduate student in 1962, and never left MIT. His MIT degrees include the S.B. in physics and electrical engineering in 1966, and the S.M. in electrical engineering in 1967. As an electrical engineering graduate student at MIT in the late 1960s and early 1970s, Pindyck focused on two areas of applied mathematics, control theory and information theory.

After having encountered and thoroughly digested Paul Samuelson's 7th edition of *Economics: An Introductory Analysis*, he switched to pursuing graduate studies in economics, completing his MIT economics department Ph.D. in 1971. His doctoral studies were supervised by MIT economics department professors Robert M. Solow and Edwin Kuh, and by MIT electrical engineering and computer science professor Michael Athans, a pioneer in the field of multivariable control systems.¹

Pindyck joined the MIT Sloan School Applied Economics faculty in 1971, and soon became engaged in the research activities of MIT's new Energy Laboratory. He has been a Research Associate at the National Bureau of Economic Research since 1983. From 2013 to 2019, he was also an Associate Scholar in the Harvard Environmental Economics Program at the Kennedy School of Government. He is an elected Fellow of the Econometric Society, has served as President of the Association of Environmental and Resource Economists, was Associate Editor of *Energy Economics*, and Co-Editor of the *Review of Economics and Statistics*. At MIT, he has received numerous outstanding teacher awards.

A large portion of Professor Pindyck's research contributions in applications of optimal control theory and analyses of behavior in dynamic markets emanated from his early MIT undergraduate and graduate studies in electrical engineering. As this is being written (March 2023), Pindyck is still at MIT, and to date he has

¹ Although Nobel laureate Robert M. Solow and econometrician Edwin Kuh are familiar names to most economists, Professor Michael Athans' name is less familiar. Professor Athans was a widely recognized pioneer in applying optimal control theory in numerous fields outside electrical engineering; for example, the National Bureau of Economic Research lists four chapters he authored or co-authored (including one with econometrician Gregory Chow), and one NBER working paper. See Professor Emeritus Michael Athans [2020] for references and further discussion.

spent 61 years there, beginning with his undergraduate enrollment in 1962; he is truly a quintessential purebred MIT alumnus.

II. INTRODUCTION TO PROFESSOR PINDYCK'S CONTRIBUTIONS

Professor Pindyck's research has been extensive, wide-ranging and rigorous. His initial research utilized optimal control theory, and focused on optimal monetary and fiscal stabilization policies, including when public authorities have conflicting objectives. Early research also examined energy pricing and the rationing that can result from price controls, the impacts of cartels and other market power industry structures, and interactions between energy policies and the macroeconomy. Subsequent research evaluated environmental policies, the social costs of deforestation and carbon emissions, and the welfare effects of catastrophes, including recent contagion phenomena such as COVID19. Among Pindyck's most notable contributions are his analyses of the impacts of uncertainty, irreversibilities, sunk costs, real options, and risk aversion on investment behavior. His research portfolio has included analyses of financial matters such as risk, inflation and the stock market, as well as the dynamics of commodity spot and futures markets. Although Pindyck's writings are uniformly grounded in economic theory, they are not just theoretical, but are often empirical and focused on implications for understanding public policies, as demonstrated in his most recent book, *Climate Future*, on the economics of climate change.

It is a formidable challenge to review such a wide and diverse range of topics in one manuscript. To keep the length of this review manageable, I confine my discussion to four areas in which Pindyck's contributions have been particularly salient: (i) economic applications of optimal control theory; (ii) market structure and the exploitation of non-renewable resource markets; (iii) uncertainty, irreversible investments, and option values; and (iv) economic issues underlying climate change and climate policy.

I begin by placing Professor Pindyck's research in a broader context, going beyond the nomenclature and notation typically used in econometric models and introducing readers to concepts and tools used in the implementation of optimal control methods.

II.A. LINEAR ECONOMETRIC MODELS

Many of the economic applications of optimal control theory utilize an econometric model as one of its components. I begin with a discussion of linear econometric models, and then turn to optimal control nomenclature.

In linear structural equation models, economic theory posits the existence of simultaneous linear relationships among individuals, sectors, or markets. Variables are typically classified as *endogenous* or *exogenous variables*, depending on the extent to which economic theory and economic behavior are intended to account for their values.²

For statistical purposes, however, a more useful distinction is between *jointly dependent variables* and *predetermined variables*. In dynamic econometric models, lagged values of the endogenous variables are distinguished from current values of the endogenous variables. The set of jointly dependent variables comprises the *current endogenous variables*; the predetermined variables consist of the exogenous variables and the lagged endogenous variables.

In traditional simultaneous equation linear econometric models, the m jointly dependent endogenous variables at time t are denoted as $y_m(t)$, the k predetermined exogenous variables as $x_k(t)$, and the unobserved random disturbances in the m^{th} structural equation as $u_m(t)$. Parameters in the m^{th} structural equation on the m'^{th} jointly dependent endogenous variable are denoted as $\gamma_{mm'}$, while parameters in the m^{th} structural equation on the k^{th} predetermined variable are denoted as β_{km} .

The t^{th} observation can be written in matrix notation as

$$y'(t)\Gamma + x'(t)\beta + u'(t) = 0. \quad (\text{Eqn. 1})$$

Eqn. 1 refers to only a single joint observation. To write this equation system in terms of all the observations, one “stacks” the T observations and writes the resulting linear simultaneous equation system in matrix notation with upper case letters as:

² A classic reference for systems of simultaneous linear equation models is the econometrics textbook by Arthur S. Goldberger [1964], especially chapter 7, “Systems of Simultaneous Linear Relationships”. For a more recent discussion with slightly different notation, see Robert S. Pindyck and Daniel L. Rubinfeld [1998], especially chapter 12, “Simultaneous Equation Estimation”.

$$Y\Gamma + X\beta + U = 0 \quad (\text{Eqn. 2})$$

where Y is a $T \times M$ matrix, X is a $T \times K$ matrix, U is a $T \times M$ matrix, and Γ and β are respectively $M \times M$ and $M \times K$ matrices as defined above. It is typically assumed that Γ is nonsingular.

In some cases, in order to facilitate computations, it is useful to transform Eqn. 2 so that only the jointly dependent variable matrix Y is on the left-hand side of each structural equation. Econometricians call this the *reduced form* transformation of the *structural equation* system Eqn. 2. Simple matrix operations yield the following reduced form specification corresponding with the structural system Eqn. 2:

$$Y = X\Pi + V \quad (\text{Eqn. 3})$$

where $\Pi = -\beta\Gamma^{-1}$, a $K \times M$ matrix, and $V = -U\Gamma^{-1}$, a $T \times M$ matrix.

Once an econometric model has been estimated, one can employ it to assess how well its dynamic properties reproduce actual historical values and trends. Other useful applications of an econometric model involve using it to simulate alternative policies, alternative trajectories of certain exogenous variables, or sensitivity to altered parameter values.

II. B. NOMENCLATURE AND NOTATION IN OPTIMAL CONTROL MODELS

In contrast to econometric models, optimal control applications attempt to identify policies that track “as closely as possible” desired trajectories of certain economic variables, but subject to tracking “as closely as possible” desired paths of other variables. Professor Pindyck’s earliest intellectual contributions involved demonstrating how optimal control techniques could be integrated with econometric models of short-run stabilization policies, thereby identifying optimal economic policies.

The variables identified in optimal control theory are typically closely related to those in econometric models. For example, it is common in optimal control applications to specify x_i as an n -dimensional vector of endogenous “state” variables, u_i as an r -dimensional vector of usually exogenous “control” variables amenable to manipulation by the policy planner, and z_i as an s -dimensional vector

of other exogenous variables that cannot be controlled by the policy planner, all at time i .

The problem facing the planner in control theory is to make x_i track “as closely as possible” a nominal, ideal or “desired” state vector x_i^* , but subject to u_i tracking “as closely as possible” to a “desired” or nominal control vector u_i^* .

In state-variable form, the discrete differential optimal control equation system of interest is written in linear form as

$$x_{i+1} - x_i = Ax_i + Bu_i + Cz_i \quad (\text{Eqn. 4})$$

with the initial condition $x_0 = \Omega$. The elements of the matrices A , B and C can be functions of the econometrically estimated parameters Γ and β in Eqn. 1 and Eqn. 3. Frequently, parameters in control theory are taken from researcher’s own or other researchers’ published research on related topics, or the authors’ perceptions of “consensus” estimates.

Econometric estimation can be used, but is not essential to the construction of an optimal control simulation model. To apply the results of an estimated econometric model to an optimal control framework, the econometric model must be respecified in state variable form such as in Eqn. 4. This yields a set of first order linear difference equations.

Suppose that x_i^* and u_i^* are the desired state and control vectors being tracked over the entire specified planning period, $i = 0, \dots, N$. At time i , the planner would like x_i to be close to x_i^* , and u_i to be close to u_i^* as possible.

To quantify how “close to” one vector is to another, calculate the x_i and u_i deviations from x_i^* and u_i^* , respectively, and then square and sum them, yielding a corresponding cost functional having the quadratic form

$$J = \frac{1}{2} \sum_{i=0}^N \{(x_i - x_i^*)'Q(x_i - x_i^*) + (u_i - u_i^*)'R(u_i - u_i^*)\} \quad (\text{Eqn. 5})$$

where Q is an $n \times n$ positive semi-definite matrix, and R is an $r \times r$ positive definite matrix.

Now the optimal control problem can be stated as finding a control sequence $\{u_i^*, i = 0, 1, \dots, N-1\}$ such that

$$x_0^* = \Omega, \quad (\text{Eqn. 6})$$

$$x_{i+1}^* - x_i^* = A_i x_i^* + B_i u_i^* + C_i z_i \quad (\text{Eqn. 7})$$

and the cost functional Eqn. 5 is minimized.

The relative magnitudes of Q and R in the cost functional Eqn. 5 quantify the costs of controlling the economy relative to having the economy deviate from its ideal or desired trajectory. Typically both Q and R are diagonal matrices, although they need not be. Some of the elements of Q may be zero. The elements of R specify the relative costs of deviating from the ideal or desired paths of the control variables --- say, the cost of deviating from the desired interest rate to the cost of deviating from the desired money supply. All of the diagonal elements of R must be non-zero, although they can be made very small and positive. While the strictly positive restriction is meaningful in terms of economic realities, it turns out also to be necessary for there to be a unique mathematical solution.

II.C. PINDYCK'S "PROOF OF CONCEPT" WITH OPTIMAL CONTROL THEORY

Based in large part on his 1971 MIT economics department Ph.D. dissertation (Pindyck [1971]) that in turn incorporated a small macroeconometric model he constructed earlier as a graduate student (Pindyck [1970]), in a series of notable publications beginning with that in the *Institute of Electrical and Electronic Engineers* professional society journal (Pindyck [1972a]), Pindyck creatively applied optimal control theory to a realistic economic policy problem -- short-run economic stabilization policy. Various aspects of his Ph.D. dissertation were subsequently published -- one in a computational economics journal (Pindyck [1972a]), another as a book in a prominent series of contributions to economic analysis (Pindyck [1973a]), and yet another in a prestigious econometrics journal (Pindyck [1973b]). Professor Pindyck was therefore among the first economists to articulate that the tools of optimal control theory (originally developed by physicists and electrical engineers) could not only in principle be applied to economic issues, but he was likely among the first actually to implement an application to a short-run economic stabilization policy issue based on an estimated macroeconometric model.³

³ References to early literature noting the possible application of optimal control theory to economic issues include Pindyck [1971; 1972a,b; and 1973a,b], and Neck [2008].

Pindyck [1972a] contains a succinct discussion of his initial application of optimal control theory to economic stabilization, which he described as a deterministic “dual discrete-time tracking problem (nominal state and nominal policy trajectories are tracked) to a linear time-invariant system with a quadratic cost functional” ([1972a],p. 287); lengthier and more detailed expositions are in Pindyck [1971; 1973a,b].

In each of these writings, Pindyck first reports results of estimating a linear quarterly econometric model of the U.S. economy over the 1955-I to 1967-IV time period, where the nine basic macroeconomic behavioral or state variables were real consumption, non-residential investment, residential investment, inventory investment, short-term interest rates, long-term interest rates, a price level, unemployment rate, and money wage rate; a tenth state variable was based on a national income and real disposable income accounting identity.⁴ The three policy or “control” variables he considered were a tax surcharge rate, government spending and quarterly change in money supply. To assess the reliability of the econometric model, Pindyck initially performed a dynamic simulation on the macroeconometric model using historic data for the policy variables. The software used in the econometric estimation and simulation was that described in Eisner and Pindyck [1973] – software at MIT that integrated econometric estimation with model solution and simulation methods.

To solve the optimal control problem analytically, Pindyck [1973a,b] expressed the necessary conditions set forth by the Pontryagin et al. [1962] Minimum Principle; because his cost functional was quadratic, his necessary conditions were also sufficient.

Pindyck began by writing the Hamiltonian based on Eqns. 5-7 above (roughly analogous to a simultaneous series of static Lagrangian minimization problems subject to a series of constraints) as

$$H(x_i, p_{i+1}, u_i) = \frac{1}{2} (x_i - x_i^*)' Q (x_i - x_i^*) + \frac{1}{2} (u_i - u_i^*)' R (u_i - u_i^*) + p_{i+1}^* (A_i x_i + B_i u_i + C_i z_i), \quad (\text{Eqn. 8})$$

⁴ Each of the equations was estimated using a two-stage least squares econometric estimation method, allowing for first-order autocorrelation using the Hildreth-Lu grid search technique. See Eisner and Pindyck [1973] for computational details.

where p_i is the vector of co-states (i.e., dynamic Lagrange multipliers or shadow values that measure the time-varying marginal cost of each state variable). He then derived the necessary conditions generating the equations describing the optimal trajectories for x_i^* , p_i^* , and u_i^* subject to boundary conditions and a transversality condition, and finally he solved the optimal control u_i^* as a deterministic linear function of x_i^* and the cumbersome but manageable matrices F_i and G_i .⁵

$$u_i^* = -F_i x_i^* + G_i \quad . \quad (\text{Eqn. 9})$$

This analytical solution assumes that at any time i , the planner knows the entire future trajectory of exogenous and endogenous variables. Although the underlying mathematics and computations may appear formidable to those not familiar with optimal control techniques, Pindyck [1973b, p. 535] claimed “the solution...really is not [formidable]. All of the above steps involve iterative solutions (and only N iterations) that require little more than multiplying and adding matrices (albeit large matrices – n might be on the order of several hundred for a large econometric model). Remember that the largest matrix that might be inverted is of dimension r , and r would normally be less than 10 and for many problems on the order of 3 or 4. On the whole, very little computer time should be required.”

Before applying the tools of optimal control theory, Pindyck needed to respecify the econometric model in state-variable form, i.e., in the form of Eqn. 4. This implied defining new state variables to replace those variables that appeared in the econometric model with lags greater than one period, and adding their definitional identities to the optimal control specification. This also involved assuming that actual values of the policy variables of government spending, tax surcharge and change in money supply in the initial time periods were the results of, and equal to, the desired levels specified at some previous point in time, such as the previous quarter. When this was done, a total of 28 state variables appeared in the Pindyck model – ten endogenous x variables and eighteen

⁵ For further details on the derivation and definition of F_i and G_i , see Eqns. 10, 11 and 12 in Pindyck [1972a], p. 289.

definitional identities involving lagged values, three u control variables, and two exogenous z variables.⁶

To demonstrate applicability of the optimal control solution technique, Pindyck [1972a,b; 1973a,b] performed several experiments (eleven in Pindyck's [1971] Ph.D. dissertation) in which the optimal control solution was used to formalize alternative stabilization policies for the econometric model. Each of the optimal policy experiments was run for 20 time periods, 1957-1 through 1962-1. However, the elements of Q and R were changed for each experiment, varying the penalties associated with alternative combinations of tax (fiscal) and monetary policies. For example, in one experiment, only the tax surcharge policy was weighted differently from other policies, whereas in another experiment the use of a tax surcharge was deemphasized while that of monetary policy was enhanced.

In all experiments the cost functional contained zero weights for those endogenous state variables considered "intermediate" in terms of policy goals. For example, the coefficients in the Q -matrix corresponding to the short- and long-run interest rates were always set to zero, although the residential and non-residential investment elements in the R matrix had non-zero coefficients. This enabled Pindyck to ignore impacts on interest rates themselves, and instead focus only on their effects on other variables of direct policy interest.

Despite its small size and simplicity, Pindyck's experimental results demonstrated "proof of concept" for optimal control and yielded several important and credible lessons for economic stabilization problems. For example, in one experiment, lengthy lags in the impact of monetary policy and in the investment sectors of the model relative to those in fiscal tax policy were incorporated, generating optimal monetary policy that was shown to be applied in short, strong bursts. The result was extreme monetary expansion at the beginning of the plan (to get the economy moving) and then monetary contraction at the end of the plan (to put the brakes on).

Pindyck summarized the findings of his optimal control policy applications by stating, "We have at least tried to show that this approach can provide a viable

⁶ A detailed presentation of the transformation of estimated econometric equations into a state-variable form of a dynamic system is in Appendix A of Pindyck's [1971] Ph.D. dissertation, pp. 243-249.

tool both for policy planning and for analyzing and better understanding a model's dynamic behavior" (Pindyck [1973b, p. 559]).

Regarding the usefulness of optimal control and optimal policy models in efficiently identifying the most preferable policy actions, Pindyck [1971, pp. 241-242] wrote: "It is important to point out that in theory there is nothing that could be learned from an optimal policy solution that could not be learned by performing simulation experiments – [if] one were willing to perform enough simulations. If one had the time and stamina to perform a large enough number of simulations, he would eventually come up with a solution, i.e., a set of paths for both the control variables and the endogenous variables, that matched the optimal policy solution. But it is exactly because this is such an inefficient approach that the optimal policy problem was posed and solved in the first place."

Pindyck's early research was grafted on to his small macro-econometric model of the U.S. economy. A subsequent research project demonstrated the usefulness of optimal control in allocating investment among five sectors in Tunisia (Martens and Pindyck [1975]), demonstrating that dynamic optimal control methods could replace then-common static linear programming tools.

In yet a different strand of research, Pindyck and Roberts [1974,1976] extended their analysis on the usefulness of control methods by examining the situation when the endogenous variable x_i vector has both "intermediate" targets such as money stock and "ultimate" targets such as GNP and price inflation. In these projects, the authors focused only on monetary policy. While tax rates and the level of government expenditures are subject to rather direct policy control, the money stock cannot be directly controlled by the Federal Reserve, although it can manipulate other variables which in turn affect the money stock. The inability to directly control these policy instruments had yielded a two-stage optimization process in which these instruments became "intermediate" targets and the true policy instruments were those variables over which the Fed had direct control, e.g., required reserve ratios, the discount rate, ceilings on interest payments on bank liabilities, and the use of open market operations to affect either unborrowed reserves or the Federal Funds rate.

In Pindyck and Roberts [1974,1976], the authors studied the optimal control problem of how a monetary authority can best manipulate the policy instruments which it can directly control in order to reach its intermediate target objectives, using a monthly money market econometric model developed at the Federal Reserve Board. In other words, they asked, what is the Fed's optimal policy given that it would like the money stock and other variables to track as closely as possible some specified time path?

Although Pindyck and Roberts [1974,1976] implemented deterministic optimal control methods as in their early research, here they also introduced a stochastic innovation: residuals from an historic simulation were used as random shocks and optimal policies were calculated by applying the deterministic control law to the model in an adaptive manner.

The primary intermediate target variables of interest were the M1 measure of the money stock and the Treasury bill rate. The principal finding from the deterministic experiments demonstrated it was very difficult to come close to the desired nominal path for the money stock, but it was not difficult to hit the interest rate exactly, suggesting that the interest rate might be a better intermediate target variable than the money stock.

These experiments were then repeated taking into account the effects of random shocks on the model, assuming the only random shocks affecting the model were additive noise terms that were independently and identically normally distributed and not autocorrelated, thereby enabling the assumptions of risk neutrality and certainty equivalence to be invoked to rationalize use of a deterministic technique.⁷

The new stochastic model, in state variable form, was specified as

$$x_{t+1} - x_t = Ax_t + B_1u_{1t} + B_2u_{2t} + Cz_t + D\varepsilon_i \quad (\text{Eqn. 10})$$

where the random error terms ε_i are generated by either adding or subtracting the residuals from a simulation of the model.

In the stochastic experiments performed in Pindyck and Roberts [1974], the presence of the random shocks did not result in a large deterioration of the

⁷ For discussion and proofs, see Theil [1957] and Chow [1972a,b;1973].

optimal control results, as long as the optimal solutions were calculated in an adaptive manner using simulation residuals.

Pindyck and Roberts summarized their findings by noting the deviations for unborrowed reserves were generally larger for the stochastic experiments than they were for the deterministic ones, but the adaptive control was self-correcting, so that the trade-off among policy instruments was not substantially worsened as long as new residual-adapted observations were used in making the next policy decision.

II.D. DECENTRALIZED OPTIMAL CONTROL WITH CONFLICTING OBJECTIVES

In Pindyck's pioneering empirical implementations of the economic stabilization optimal control policy described above, he assumed a single central planning authority attempted optimally to control government spending, changes in tax policy, and money supply. How do the theory of optimal control, the resulting optimal control policies, and values of the state variables change when instead of a single central planner, there are two planning authorities, who may or may not have conflicting objectives? That is the focus of Pindyck's [1976,1977] next stage of research.

In the U.S., monetary and fiscal policies are exercised by separate authorities that are largely independent of each other, and that may have conflicting objectives. This separation of monetary and fiscal control may considerably limit the ability of either authority to stabilize the economy, particularly when the conflict over objectives is at all significant. Moreover, because monetary policy operates with longer lags than fiscal policy, the proper time-phasing of the two can be critical.

In Pindyck [1976], it is assumed that each authority arrives at its policy using the same econometric model (i.e., each has the same view of the way the world works), but that the two have different sets of objectives. The econometric model is linear and deterministic, so that it can be represented in state variable form as

$$x_{t+1} - x_t = Ax_t + B_1u_{1t} + B_2u_{2t} + Cz_t \quad (\text{Eqn. 11})$$

with initial condition $x_0 = \Omega$. Here x_t is a vector of n state variables, u_{1t} and u_{2t} are vectors of r_1 and r_2 policy (control) variables manipulated by the fiscal and monetary authorities, respectively, and z_t is a vector of s uncontrollable

exogenous variables whose future values are known and can be predicted with certainty. A , B_1 , B_2 and C are $n \times n$, $n \times r_1$, $n \times r_2$, and $n \times s$ matrices, respectively.

Each authority chooses an optimal trajectory (a “strategy”) for its own set of control variables over the time period $t = 0, 1, \dots, N-1$. The first authority chooses its strategy $\{u_{1t}\}$ to minimize its cost functional

$$J_1 = \frac{1}{2} (x_N - x_{1N}^*)' Q_1 (x_N - x_{1N}^*) + \frac{1}{2} \sum_{t=0}^{N-1} \{ (x_t - x_{1t}^*)' Q_1 (x_t - x_{1t}^*) + (u_{1t} - u_{1t}^*)' R_{11} (u_{1t} - u_{1t}^*) + (u_{2t} - u_{2t}^*)' R_{12} (u_{2t} - u_{2t}^*) \} \quad (\text{Eqn. 12})$$

and the second authority chooses its strategy $\{u_{2t}\}$ to minimize its cost functional

$$J_2 = \frac{1}{2} (x_N - x_{2N}^*)' Q_2 (x_N - x_{2N}^*) + \frac{1}{2} \sum_{t=0}^{N-1} \{ (x_t - x_{2t}^*)' Q_2 (x_t - x_{2t}^*) + (u_{1t} - u_{1t}^*)' R_{21} (u_{1t} - u_{1t}^*) + (u_{2t} - u_{2t}^*)' R_{22} (u_{2t} - u_{2t}^*) \}. \quad (\text{Eqn. 13})$$

Here x_{1t}^* and x_{2t}^* represent nominal or “desired” values for the state variables from the points of view of authorities 1 and 2, respectively, and similarly u_{1t}^* and u_{2t}^* represent desired values of the control variables for each authority. The matrices Q_1 and Q_2 represent for each authority, the relative weights assigned to deviations from the desired paths for each state variable. R_{11} and R_{22} designate the relative weights that each authority assigns to deviations from the desired path for *its own* control variables, while R_{12} and R_{21} designate the relative weights that each authority assigns to deviations from the desired path for the other’s control variables; thus, R_{12} and R_{21} indicate how important it is for each authority that the other authority stay close to its policy variable targets

Pindyck then borrows from the engineering and physics literature the concepts of *open loop* and *closed loop* strategies. Specifically, when the centralized planning agency designs its optimal policy at the beginning of the planning period, given its objectives as specified in the cost functional and the initial conditions, and then adheres to that policy throughout the entire planning period, this is called an *open-loop* strategy. In effect, the optimal control u_i^* depends, at any time i , on the initial condition x_0 . In contrast, when the planning authority designs a control role at the beginning of the planning period, and then uses that control role, together with observations on realizations of the state of the economy, to continuously revise and adapt its control policy, this is called a *closed-loop* strategy. In this latter case the optimal control u_t^* depends on the current state x_t .

The optimal control policy is essentially a discrete-time differential game. Pindyck defines the Nash solution to the (u_1^*, u_2^*) game as satisfying the conditions

$$J_1(u_1^*, u_2^*) \leq J_1(u_1, u_2^*) \quad (\text{Eqn. 14})$$

and

$$J_2(u_1^*, u_2^*) \leq J_2(u_1^*, u_2) \quad (\text{Eqn. 15})$$

for all possible u_1 and u_2 . What Eqns. 14 and 15 require is that for the first authority, its value of J_1 at u_1^* , given authority 2 is at u_2^* , must be at least as small as J_1 evaluated at any other values for u_1 , and simultaneously, for the second authority, its value of J_2 at u_2^* , given authority 1 is at u_1^* , must be at least as small as J_2 evaluated at any other values for u_2 .

Pindyck [1977, pp. 520-523] analytically derived the deterministic solution to both the open-loop and closed-loop control problems. He also generalized to the case of k controllers, rather than just two

Pindyck [1976,1977] applied the open-loop solutions to the same small econometric model examined in his previous [1971, 1972a,b, and 1973a,b publications]. Fiscal policy was incorporated through exogenous government expenditures and a surtax, and monetary policy through the money supply (currency plus demand deposits). Both authorities were assumed to agree on what is the desired unemployment rate and rate of price inflation, but fiscal authorities placed greater weight on unemployment while the monetary authority placed greater importance on inflation (as indicated in elements of their diagonal Q matrices). Desired or nominal trajectories were the same as those used in Pindyck's previous publications – the fiscal authority desires a zero surtax and steady growth in government expenditures, and the monetary authority desires steady growth in the money supply. The fiscal authority can manipulate only government spending, and the monetary authority only the money supply.

Results indicated it is considerably easier to reduce the unemployment rate than it is to reduce the inflation rate, so that when both were weighted equally in the cost functionals the optimal policy favored the unemployment rate. In one experimental run in Pindyck [1977] each authority had only one target variable – the price level for the fiscal authority and the unemployment rate for the

monetary authority, but these conflicting objectives were reversed in another run. A comparison of runs indicated that the fiscal authority had more control over the economy than did the monetary authority, reflecting the facts that in the common underlying econometric model the fiscal multipliers were larger than the monetary multipliers, and that longer time lags were inherent in monetary policy.

Pindyck noted, however, that this apparent superiority of fiscal policy reflected some controversial and possibly inaccurate assumptions in his modeling: While monetary policy operated with longer lags than fiscal policy, in fact fiscal variables (government spending and tax policy) cannot be manipulated as frequently and as freely as monetary variables. This limitation on the ability of the fiscal policy to manipulate its variables would probably reduce significantly the fiscal “advantage” that Pindyck observed. However, Pindyck’s results seemed to indicate that the suboptimality resulting from the conflicting objectives of the two authorities becomes severe only in the first four to six quarters of the planning period. After six quarters a “compromise” behavior occurred where neither authority was as close to its desired targets as it would be in a cooperative situation, yet the deviations from targets were not exacerbated in later years of the planning period, and oscillating outcomes did not occur.⁸

Pindyck recognized that the linear-quadratic optimal control specification is not constrained to yield numerically “sensible” values. For example, in one of Pindyck’s experimental runs, the optimal solution yielded a negative unemployment rate (Pindyck [1971], Run 8). Another limitation of the linear-quadratic specification is that it is restrictive in that the cost functional is symmetrical, i.e., overshooting a policy target incurs the same cost as undershooting the target.

Pindyck acknowledged there may be functional forms other than the quadratic that are more representative of actual social costs, but reliably parameterizing them would be a very complex task. Attempts to generalize the cost-functional include a piece-wise quadratic, or a conjugate gradient method. As noted in Pindyck [1972b], even in relatively small nonlinear models, it has been difficult to obtain numerical convergence for solutions of optimal control. He noted in Pindyck [1972b] that “The experience in engineering has been that often the

⁸ More nuanced behavioral assumptions by the two authorities are considered in Neese and Pindyck [1984].

closed-loop control for a linear model can be applied adequately to the control of a physical system that is nonlinear. We have had less experience with control theory in economics, but we can expect that the adequacy or inadequacy of linear or linearized models will depend on how much of the dynamic behavior of the economic system is determined by the nonlinearities in its structure. Our analytical tools for dealing with the dynamics of nonlinear systems are meager and so we may have to look at computational results to get a better feeling for how much we can rely on linear optimal control as a means of obtaining stabilization policies.”

Notably, in his [1972b] article, Pindyck advocated using econometric simulation methods extensively, rather than pursuing nonlinear optimal control methods, to identify preferable time paths for various policies. Recognizing today that Pindyck was writing in 1972, more than half a century ago when many computational advances had yet to be invented, Pindyck asserted the importance of computational feasibility criteria that in his view favored simple optimal control techniques: “If optimal control is ever to gain the acceptance that simulation has as a practical tool for policy planning and analysis, it is imperative that it yield solutions that are computationally tractable...The linear-quadratic specification is robust in its applicability to the stabilization problem, and has the special advantage of being computationally tractable” (Pindyck [1972b, p. 389]).

III. CARTELS AND NATURAL RESOURCE EXPLOITATION

Robert Pindyck had a strong interest in global energy and natural resource markets including various market structures, such as monopolies and cartels, primarily in natural resource industries. In analyzing these markets, Pindyck used the tools of game theory.

III.A GAME THEORY: OPEC AS A SUSTAINABLE CARTEL

Hnyilicza and Pindyck [1976], for example, highlight that the composition of the OPEC cartel was diverse with potentially conflicting objectives: OPEC contained a block of “spender” countries with large cash needs and relatively small proven resource reserves (e.g., Iran, Venezuela, Indonesia, Nigeria, Algeria, and Ecuador), and another block of “saver” countries with little need for cash, a lower discount rate, and a relatively large proven resource reserve base e.g., Saudi Arabia, Abu Dhabi, Bahrain, Kuwait, and Qatar. If the two groups had the same objectives,

they could act together as a monopolist. But if they act as a cartel with different desired policies, actual cartel policy would depend on an agreement between the two groups accommodating both differences in objectives and in bargaining power. What would cartel policy look like with two such diverse sets of players?

The approach Hnyilicza and Pindyck [1976] took is based on the theory of *cooperative (not non-cooperative) games* developed by John Nash [1953]. Hnyilicza and Pindyck derived optimal trajectories for both price and the ratio of output shares assuming that the cartel maximized a weighted sum of discounted profits for each of the two groups of countries.

The two-part cartel consisted of a group of saver countries that have the objective

$$\text{Max } W_1 = \sum_{t=1}^N (1/(1+\delta_1)^t) [P_t - m_1/R_t^1] D_t^1 \quad (\text{Eqn. 16})$$

and a group of spender countries with the objective

$$(\text{Max } W_2 = \sum_{t=1}^N (1/(1+\delta_2)^t) [P_t - m_2/R_t^2] D_t^2 \quad (\text{Eqn. 17})$$

where P is the price per barrel of oil in real terms, R^1 and R^2 are oil reserves in country groups 1 and 2 with $R^1 > R^2$, and δ_1 and δ_2 are the discount rates for country groups 1 and 2, respectively with $\delta_1 < \delta_2$. The ratio variables m_1/R_t^1 and m_2/R_t^2 are the average production costs (so that parameters m_1 and m_2 determine initial average costs) for each group of countries. D_t^1 and D_t^2 are the production of each group of countries determined by a division of total cartel production

$$D_t^1 = \beta_t D_t, \quad (\text{Eqn. 18a})$$

$$D_t^2 = (1 - \beta_t) D_t \quad (\text{Eqn. 18b})$$

where D_t is total global demand for cartel oil in billions of barrels of oil per year, with $0 \leq \beta_t \leq 1$.

Suppose a cooperative agreement is worked out whereby price and output shares are set to maximize a weighted sum of the objectives of each group of countries

$$\text{Max } W = \alpha W_1 + (1 - \alpha) W_2, \quad 0 \leq \alpha, \beta_t \leq 1. \quad (\text{Eqn. 19})$$

There are now two control variables for the cartel – price and output share, and both can vary over time. By varying α between 0 and 1 and solving the resulting set of parametric optimization problems, a Pareto-optimal frontier emerges in the

space of realized outcomes (W_1, W_2). Each point on the frontier corresponds to a different trade-off between the relative objectives of the two groups of countries.

Determining the value of α that is most likely to prevail as a result of a negotiated agreement between the two groups of countries requires the solution of a cooperative two-person game, i.e., requires a theory of bargaining. Hnyilicza and Pindyck [1976] utilize an extremely general and robust theory of bargaining put forth by John Nash [1953], two years after he developed his better known non-cooperative theory. (See Figures 2 and 3 in Hnyilicza and Pindyck [1976, pp. 145-146] for a graphical presentation.)

Since each of the two parties in a bargaining game attempt to move along the set of bargaining outcomes in opposite directions, the problem is to determine a meaningful measure of market power for the two parties. In the Nash approach, the notion of a *threat point* is introduced, i.e. the outcome that would result if negotiations were to break down and non-cooperative behavior were to ensue. In essence, Nash's solution is based on the premise that the relevant measure of 'relative power' that determines the outcome of the bargaining process is given by the relative utilities at the status quo or at the point of no agreement. This is plausible, since the reason each party is willing to bargain is that it expects to achieve a payoff over and above the payoff attained at the threat point – both parties should be willing to accept a division of the net incremental gains in a proportion directly related to the losses incurred by not making an agreement.

As is common in cooperative games, for any value of α , the optimal path for β_t follows a 'bang-bang' solution. In particular, the optimal β_t will remain at zero for some time (until spender country reserves are depleted) and then jump to 1 (where it will remain until saver country reserves are depleted). However, this cooperative solution may not be politically feasible; because of the incentive to cheat, it may not be realistic to expect the two groups of countries to agree to this allocation of output. Instead one might expect the two groups to divide output in proportion to historic production levels, and simply optimize with respect to a single price that approximates the monopoly price, with little need to negotiate. On the other hand, if the output shares reflect choices and negotiations, the optimal paths will depend significantly on the relative bargaining power of each block.

Hnyilicza and Pindyck then argue that a “compromise” policy might be adopted whereby saver countries initially cut back production more than spender countries, but then expand production after 10 or 15 years, either with agreed-upon cutbacks by spender countries (who by then may have exhausted a significant fraction of their reserves) or with a drop in price. They conclude, “In fact, such a compromise policy may be exactly what we are observing now” (Hnyilicza and Pindyck [1976], p. 153). Note these authors were writing in 1976, not 2023, and in the half century since they wrote there have been other seismic shifts in oil markets. One prediction they made is clearly accurate, however: “Recognizing that the cartel consists of producers with somewhat different interests will be essential in predicting its response to these future cutbacks” (p. 153).

III.B. WHAT MAKES CARTELS FEASIBLE AND SUSTAINABLE?

OPEC is not the first cartel for an exhaustible resource, nor is it likely to be the last. In Pindyck [1978a], the issue addressed is “what makes a natural resource cartel feasible and sustainable”? A static analysis is unlikely to be informative, he argues, for the process of reserve depletion might have an important impact on monopoly pricing decisions, and on the potential gains from cartelization. Moreover, to the extent demands and supplies adjust only slowly to changes in price, a cartel might have the potential for large short-term monopoly profits by taking advantage of adjustment lags, behavior that is ignored in static analyses.

Pindyck [1978b] treats a cartel as a pure monopolist holding a known quantity of reserves and facing a ‘net demand’ function (total world demand minus supply from ‘competitive fringe’ producers who are not members of the cartel). He reasons that if the gains to the pure monopolist are small, one should not expect the cartel to remain stable over a long time period, while if the gains to cartelization are quite large, there should be sufficient incentive for the producers to overcome the problems typical of cartelization.

Pindyck then examines three known cartels: OPEC for oil, CIPEC (International Council of Copper Exporting Countries) in the case of copper, and IBA (International Bauxite Association) in the case of bauxite. The market dominance of these cartels varied considerably in the 1970s: while OPEC and IBA accounted for about two-thirds of non-Communist world oil and bauxite production, CIPEC

accounted for only one-third of copper production. For each market, Pindyck computes the optimal price trajectory and optimal sum of discounted profits for a monopolist cartel, and then compares these with the optimal price trajectory and sum of discounted profits that would result if the cartel dissolved (or never formed) and its members behaved competitively.

Because of adjustment lags, it was optimal for OPEC to charge a high price initially, taking advantage of the fact that net demand can only change slowly. He finds the optimal monopoly price rises slowly to about \$15 a barrel in 1975, declines to about \$10 for the next five years, and then rises slowly. The relative gains from cartelization by OPEC are largest during the first five years, since it is during these years the monopoly cartel can take advantage of adjustment lags and reap large short-term profits. Pindyck concludes the gains to OPEC from cartelization were high under a broad range of assumptions.

With bauxite, for a range of prices up to about \$15.60 per ton the demand for bauxite is quite inelastic, but at higher prices it becomes economical to produce aluminum from sources other than bauxite, so that the demand for bauxite becomes almost infinitely elastic. In the inelastic region, the demand for bauxite depends on the demand for aluminum, with short-run elasticities much smaller than long-run elasticities (in absolute value). At a price of about \$15.60, Pindyck expects the demand for bauxite to fall rapidly to zero, i.e., there is a “limit price” for bauxite of about \$15.60.

Bauxite is quite abundant; reserves for the competitive fringe could sustain production for nearly 300 years at the then current levels. Initial production costs were about \$5 per ton. In both the monopoly and competitive cases depletion plays only a very small role in price determination for the first 30 years; competitive markups over costs are quite small. However, monopoly price can almost be chosen at the profit-maximizing “limit” price each period, ignoring future periods. By comparison, the competitive price is lower than the monopoly price initially, but rises very slowly. The gains from cartelization are large, relatively larger for bauxite than for oil. The large gains, Pindyck concludes, should be sufficient for the maintenance of the IBA cartel.

Copper is quite different from oil and bauxite, primarily because of the large secondary supply, i.e. production from scrap copper. The presence of secondary

copper from scrap limits adjustment lags, for both monopoly and competitive producers. Unlike demand and competitive supply of oil and bauxite that adjust only slowly to changes in price allowing large short-term gains to a cartel, secondary copper supply responds very rapidly, thereby limiting gains to cartelization. To Pindyck [1978a], it is not surprising that it has been challenging for countries to establish and sustain a copper cartel.

III.C. CHALLENGING HOTELLING ON THE PRICING OF NON-RENEWABLE RESOURCES

The distinct trajectories observed and analyzed by Pindyck [1978a] raise the issue of what is the general theory of pricing for exhaustible resources. The classic theory is that from Hotelling [1931], who first demonstrated that if extraction costs were constant, under competition price minus marginal cost should rise at the rate of discount r , while in a monopolistic market, rents (defined as marginal revenue minus marginal cost) should rise at the rate of r .

However, this simple version of Hotelling's r -percent growth rate for the price of an exhaustible resource has had only limited success in reproducing the actual evolution of resource prices; the prices of most exhaustible resources have not risen steadily over time, but instead have experienced long secular declines, or more commonly, have fallen over a long period and then later risen, following a U-shaped profile. Pindyck [1978b] and Levhari and Pindyck [1981] provide theoretical frameworks that can rationalize not only the Hotelling r -percent growth scenario, but also the much more frequently observed U-shaped price history. What is this general theory?

Pindyck [1978b] begins by arguing for replacing the word "exhaustible" with "non-renewable", since the concepts of reserves and their exhaustion are ultimately economic rather than geological or physical notions.

Rather than a finite fixed set of reserves being depleted over time, Pindyck envisages reserves as being created and depleted by producers facing economic incentives. Producers are not "endowed" with reserves but instead must develop them through the process of exploration. In Pindyck [1978b], reserves are increased or maintained via exploratory behavior. Exploratory activity is the means by which reserve levels are accumulated or maintained, and depletion is treated by assuming that reserve additions ("discoveries") resulting from

exploratory activity fall as cumulative discoveries increase. Pindyck points out that “potential reserves” are unlimited, but as depletion ensues, repeated amounts of exploratory activity result in ever smaller discoveries. Given these constraints, resource producers must simultaneously determine their optimal rates of exploratory activity and production/depletion.

In Pindyck [1978b], the optimal intertemporal exploratory-production strategy is derived for competitive and monopolistic markets. He shows that if the initial reserve endowment is small, the price profile will be U-shaped; at first production will increase as reserves are developed, and later production will decline as both exploratory activity and the discovery rate fall.

Exploratory activity is chosen to build the reserve base up to a level that reduces extraction costs and then is adjusted over time so as to trade off cost savings from postponed and discounted exploration with savings from lower extraction costs and revenue gains from greater total production. Thus, the pattern of optimal exploratory activity depends highly on initial reserve levels and on rates of depletion.

If the initial reserve endowment is small, the price profile will be U-shaped, rather than steadily increasing as in the Hotelling model. This helps explain the fact that real prices of many non-renewable resources have fallen over the years. For example, the decline of real oil prices prior to the formulation of OPEC, and the decline in the real price of bauxite prior to the cartelization of the world bauxite market, can be attributed to the significant increases in the proven reserves of those resources that allowed production to increase steadily.

In the later stages of resource use (or throughout, if the initial resource endowment is large) price will increase over time as in the Hotelling framework. However, the introduction of exploratory activity reduces the rate of increase in price. Finally, in the development of a new resource for which depletion is not significant (but for which exploration and reserve accumulation are necessary), an optimal steady-state reserve level is reached that is independent of any initial reserve endowment.

Levhari and Pindyck [1981] have a different focus from Pindyck [1978b]. They observe that many resources are *durable*, so that their demands are for quantities of *stock in circulation*, rather than for *flows of production*, and thus depend on

expected *changes in prices* (“capital gains” or “capital losses”) as well as the current price level. As examples of durable resources, Pindyck-Levhari [1981, p. 366] cite “diamonds, gold, silver, and the other precious metals, but copper and other non-precious metals also have durable aspects”. Levhari and Pindyck also allow for *partially durable resources*, i.e., resources that depreciate over time, and distinguish them from *perfectly durable goods*.

Levhari and Pindyck [1981, p. 367] go on to state that “With nondurable resources such as oil and gas, demand is a flow, since once a unit of the resource has been consumed (burned), it no longer provides utility. A unit of a durable resource, on the other hand, continues to provide utility as long as it is held, so that demand is a stock relationship.”

Levhari and Pindyck [1981] show that Hotelling’s r -percent rule applies to a partially or totally durable resource produced in a competitive market, but does not apply if the resource stock is produced in a monopolistic market. However, they stress the r -percent rule clearly does not mean that price is steadily rising. Rather, if marginal production cost rises with the rate of production (which, they argue, must for the problem to make much sense), the competitive market price will fall initially as the stock in circulation increases, and later will rise as the stock decreases and eventually depreciates asymptotically to zero after production has ceased. In the special (and rare) case the resource is perfectly durable and demand is static, price will always be falling, and if the resource is partially durable but demand is growing at a rate close to (but less than) r , it is possible for price always to be rising. Levhari and Pindyck argue these are very special and extreme cases, and that in general one would expect the price of a durable exhaustible resource (a term they continue to use in spite of Pindyck [1978b]) to follow a broad U-shaped trajectory.

Levhari and Pindyck [1981] conclude, therefore, that while there are many factors and scenarios that can explain the historically observed U-shaped trajectory of exhaustible resources, for some resources durability may play a particularly prominent role in explaining observed price behavior.

III.D. MASSIVE UNCERTAINTY IN NON-RENEWABLE RESOURCE MARKETS

After having devoted considerably theoretical effort to helping explain why observed price behavior of non-renewable resources have historically followed a

U-shaped trajectory rather than obeying Hotelling's r -percent rule, in a series of subsequent research projects Pindyck focused attention on the role of uncertainty – lack of knowledge regarding what demand and reserves will be in the future – in affecting market price evolution, the optimality of competitive markets, and the role and value of exploration in affecting uncertainty.

Pindyck's [1980] characterization of uncertainty is not the common phenomenon in which some parameter or variable is taken to be unknown. He models demand uncertainty by assuming that the market demand function shifts randomly but continuously through time according to a specific stochastic process with independent increments. Although today's demand is known precisely, future demand may be larger or smaller and has a variance that increases with the time horizon. Similarly, Pindyck envisages reserve uncertainty assuming that available reserves shift upward or downward, again randomly.

Hence, as exploitation proceeds, resource producers may find that more or less reserves are available for production than originally anticipated. In such a world, the observed market price will be a random process; Pindyck [1980, 1984] addresses a number of questions about the behavior of the market in expected value terms: (i) Does the presence of such uncertainty affect the behavior of the market, e.g., should the presence of uncertainty cause producers in competitive or monopolistic markets to be more or less "conservationist" than if there were no uncertainty? (ii) Do competitive markets exploit the resource at a rate that is socially optimal in the presence of uncertainty? And (iii) what are the implications of uncertainty for exploration, either as a means of reducing the uncertainty itself, or simply to accumulate reserves?

Producers in Pindyck's [1984] model have complete information about the current status of a renewable resource market; what they do not know is the values of demand, and what the level of reserves will be in the future. However, since stochastic fluctuations occur continuously over time, producers can adapt to those fluctuations continuously. As a result, stochastic fluctuations alter the expected rate of change of price or resource extraction only to the extent that the average cost of production or productivity of exploration is changed through nonlinearity in a fluctuating variable.

Thus with average production cost constant, price will rise according to Hotelling's r -percent rule. However, even with average production cost constant, the rate at which production falls, and the initial values of production and price, are affected by uncertainty. This occurs for two reasons. First, fluctuations cause shifts in demand functions, and second, because if demand is nonlinear, zero-mean fluctuations in price imply a net change in production in order for markets to clear.

Regarding exploration, as a means of gathering information, i.e. to reduce the variance of stochastic reserve fluctuations, exploration should only be used if production costs vary with reserves. In particular, if average costs decline with production, knowledge of reserves over time permit production costs to be reduced on average by allocating more production to periods when reserves are known to be larger. When instead exploration is used to accumulate reserves, the time profile of exploratory activity is altered if a stochastically fluctuating parameter enters the discoveries function nonlinearly. This occurs because fluctuations can change the average productivity of exploratory effort and thereby shift the optimal level of exploration.

III.E. MEASURING MARKET POWER IN DYNAMIC MARKETS

Pindyck's contributions to the economics of non-renewable resources typically emphasize the role of market structure in affecting output quantity and price. But how does one measure market power when markets are dynamic, i.e., when price and output are determined intertemporally and jointly? That is the focus in Pindyck [1985].

The well-known measure of monopoly power is the Lerner Index, $L = (P - MC)/P$. In competitive markets, $P = MC$ and $L = 0$. In a static market, for a monopolist $L = 1/\eta_f$, where η_f is the absolute value of the price elasticity of demand facing the firm, implying that the firm's elasticity of demand completely determines its market power. The larger is L , the greater is the degree of monopoly power.⁹ Pindyck [1985] argues persuasively that in dynamic markets, L is a misleading measure of monopoly power, in some cases understating monopoly power, and in other cases overstating it. In Pindyck [1985], using several examples, the

⁹ For an introductory discussion, see Pindyck and Rubinfeld [2013], pp. 371-373.

misleading nature of L as a measure of monopoly power is discussed: non-renewable resource markets, markets in which firms encounter learning curves, markets where firms face adjustment costs for quasi-fixed factors of production, and markets where firms' demands respond over time (rather than instantaneously) to changes in price. In each example, L applies to only an instant time, while the impact of monopoly power always applies to some interval of time.

Pindyck's generalization of L simply involves defining marginal cost more fully so that it incorporates any relevant "user costs" – the sum of discounted future costs or benefits that result from current production decisions, i.e., intertemporal externalities. He also considers measurement of monopsony power.

Pindyck argues that the Lerner index should be altered as follows

$$L^*(t) = (P_t - FMC_t)/P_t = 1 - (FMC_t/P_t) \quad (\text{Eqn. 20})$$

where FMC is the full marginal cost at time t , evaluated at the monopoly output level. By this, Pindyck means that FMC_t includes any positive or negative what he calls "user costs" that result from the intertemporal nature of the firms' optimization problem, where these user costs are calculated under the assumption that the firm is competitive. With FMC calculated in this way, $0 \leq L^*(t) < 1$ for all t , and $L^*(t) = 0$ in a perfectly competitive market.

Consider, for example, a market in which producers learn from producing, so that average and marginal costs depend in part on cumulative production or "learning by doing"; in this case, with additional current production reducing future average and marginal costs, current marginal revenue at current production levels will exceed current marginal costs by a "user cost" amount representing the change in the discounted producer plus consumer surplus due to the value to the monopolist of producing one more unit of cumulative output. In this case, Pindyck shows that use of the traditional Lerner Index $L(t)$ would understate dynamic monopoly power as measured by $L^*(t)$.

By contrast, in the case of a nonrenewable resource, to the extent additional current production from a fixed reserve base increases future extraction and production costs, $MC(t) > FMC(t)$, and $L(t) > L^*(t)$.

A third example of monopoly power mismeasurement by the traditional Lerner Index involves dynamic demand functions, when consumers adjust their spending patterns, or as other competitive firms expand their production capacity as prices rise, implying that short- and long-run price elasticities differ, with the latter being larger in absolute value.¹⁰ Suppose, for example, the monopolist's demand function is more elastic in the long run than in the short run. Then it is optimal for the monopolist initially to set output above the point where marginal cost equals short-run marginal revenue; doing so creates a future marginal benefit by retarding the response of demand and the adjustment to long-run equilibrium. In this case, monopoly power measured by $L(t)$ is less than that by $L^*(t)$. Just the opposite occurs if the monopolist's demand curve is more elastic in the short- than in the long run, as occurs when demand evolves according to "stock-adjustment" behavior, e.g., as occurs with copper where the durable good has a secondary or scrap supply.

IV. UNCERTAINTY, IRREVERSIBLE INVESTMENTS, AND OPTION VALUES

In the 1980s, Robert Pindyck was among a number of economists who expanded our understanding of investment behavior by considering implications of the fact that most investments involved irreversible decisions, leading to opportunity costs.¹¹

IV.A. IRREVERSIBILITIES MAKE NET PRESENT VALUE CRITERIA UNRELIABLE

Pindyck [1988, p. 969] begins by asserting most investments involve some degree of irreversibility, and arguing an implication is that the Net Present Value rule of investment behavior is often misleading: "Most major investment expenditures are at least partially irreversible: the firm cannot disinvest, so the expenditures are sunk costs. Irreversibility usually arises because capital is industry- or firm-

¹⁰ A fourth example of monopoly power mismeasurement is discussed in Pindyck [1986], in which internal costs of adjustment are considered for quasi-fixed factors of adjustment. Adjustment cost models are estimated in Pindyck and Rotemberg [1983a,b].

¹¹ Pindyck's contributions to this literature include Majd and Pindyck [1987], Pindyck [1988], He and Pindyck [1992] and Pindyck [1993a,b]. Here we focus primarily on Pindyck [1988]. Majd and Pindyck focus on the time between making a decision to invest vs. the time it takes to complete the investment project ("time to build") during which the firm faces irreversible options on whether and how rapidly to complete the project. A principal finding is that increased uncertainty in demand likely has a depressive effect on the level of investment, an effect that becomes larger when there is time to build. He and Pindyck consider irreversible investment in the context of a multi-output firm that can install output-specific capital or, at greater cost, flexible capital that can be used to produce different outputs. A principal finding is that flexible capital is preferred only if its cost premium is low.

specific, that is, it cannot be used in a different industry or by a different firm. A steel plant, for example, is industry-specific. It can only be used to produce steel, so if demand for steel falls, the market value of the plant will fall. Although the plant could be sold to another steel company, there is likely to be little gain from doing so, so the investment in the plant must be viewed as a sunk cost. As another example, most investments in marketing and advertising are firm-specific, and so are likely sunk costs....When investment is irreversible and future demand or cost conditions are uncertain, an investment expenditure involves the exercising or 'killing' of an option – the option to productively invest at any time in the future. One gives up the possibility of waiting for new information that might affect the desirability or timing of the expenditure; one cannot disinvest should market conditions change adversely. This lost option value must be included as part of the investment. As a result, the Net Present Value (NPV) rule 'Invest when the value of a unit of capital is at least as large as the purchase and installation cost of the unit' is not valid. Instead, the value of the unit must *exceed* the purchase and installation cost, by an amount equal to the value of keeping the firm's option to invest these resources elsewhere alive – an opportunity cost of investing." Pindyck [1988, p. 969] then cites previous work in Majd and Pindyck [1987] showing that in many cases projects should be undertaken only when their present value is at least double their direct cost.

Pindyck [1988] shows that a firm's capacity choice is optimal when the present value of the expected cash flow from a marginal unit of capacity just equals the total cost of that unit, the latter including the purchase and installation cost plus the opportunity cost of exercising the option to buy the unit. A firm's market value has two components: the value of installed capacity (i.e., the value of the firm's options to utilize some or all of this capacity over time), and the value of the firm's options to add more capacity later. Pindyck's numerical calculations suggest that for many firms, 'growth options' should account for a substantial fraction of market value, and the more volatile is demand, the larger is this fraction. In many investment projects, incremental investment occurs sequentially, with the firm investing until the value of a marginal unit of capital is equal to its total cost – the purchase and installation cost, plus the opportunity cost of irreversibly exercising the option to invest in the unit, rather than waiting for more information and keeping the option alive.

From where do such firm-specific productive investment opportunities arise? In some cases it is the result of a patent on a production technology, or ownership of land or natural resources. More generally, a firm's managerial resources, reputation, market position and possibly, scale, all of which may have been built up over time, enable it to productively undertake investments that individuals or other firms cannot undertake (Pindyck [1988, p. 970]).

To the extent firms rely on the NPV rule in their investment decisions, Pindyck concludes, they will tend to overinvest. In markets with volatile and unpredictable demand, firms should hold less capacity than they would if investment were reversible, or future demands were known. Pindyck emphasizes much of the market value of firms is due to the possibility (as opposed to the expectation) of increased demands in the future. This helps explain that investment often occurs in spurts, and only when demand is rising above historic levels (Pindyck [1988], p. 980).¹²

In concluding, Pindyck [1988, p. 983] inquires whether firms correctly take into account the opportunity cost of investing when making expansion decisions. Citing empirical evidence, he finds that for manufacturing firms, market values tend to increase (decrease) when managers announce an increase (decrease) in planned investment expenditures, which is inconsistent with a systematic tendency to overinvest. But he then cites anecdotal evidence that managers often base investment decisions on present values computed with discount rates that far exceed those that would be implied by the Capital Asset Pricing Model – diversifiable and nondiversifiable risk are sometimes confused, and an arbitrary “risk factor” is often added to the discount rate. He concludes, “It may be, then, that managers use the wrong method to get close to the right answer” (Pindyck [1988], p. 983).

IV.B. COST VS. DEMAND UNCERTAINTY AND IRREVERSIBLE INVESTMENTS

The above discussion regarding Robert Pindyck's contributions to the literature on uncertainty and irreversible investment focuses primarily on the effects of *demand* uncertainty and future payoffs. Sometimes the *cost* of an investment is more important than the future payoff, particularly for large projects that take considerable time to build, such as nuclear power plants, large petrochemical

¹² Also see Caballer0 and Pindyck [1996]

complexes, new lines of aircraft design, and urban construction projects, although large size is not a requisite.

Pindyck [1993b] examines implications of cost uncertainty for irreversible investments. Two types of uncertainty emerge when the time to complete an investment projects takes considerable amount of time. The first is *technical uncertainty*, and it relates to the physical difficulty of completing a project; while the prices of construction inputs are known, the costs, time, effort, and materials required are uncertain. Technical uncertainty can only be resolved by undertaking the project; actual costs and construction time unfold as the project proceeds, regardless of whether total construction costs are less or greater than anticipated. Moreover, technical uncertainty is largely diversifiable. It results only from the inability to predict how difficult a project will be, which is likely to be independent of the total economy.

The second type of uncertainty relates to *input costs*, and is external to what the firm does, arising when the prices of labor, land and materials needed to build a project fluctuate unpredictably, or when unpredictable changes in government regulations change construction costs. Prices and regulations are exogenous, occurring whether or not the firm is investing, and are more uncertain the farther into the future one looks. Input cost uncertainty is particularly important for projects that take time to complete or are subject to voluntary or involuntary delays. Input cost uncertainty may be partly nondiversifiable, since changes in construction costs are likely to be correlated with overall economic activity.

The sources and amounts of cost uncertainty will vary greatly across different projects. Pindyck argues, however, that based on the range of parameter values that would apply to the bulk of large capital investments, factor cost uncertainty is likely to be more important than technical uncertainty in terms of its effect on the investment rule and the value of the investment opportunity. Pindyck [1993b] demonstrates this for investments in nuclear power plants, but he notes the opposite may be the case for some biopharmaceutical R&D projects. Although Pindyck finds that the critical cost to completion is not very sensitive to the degree of technical uncertainty, this finding is based on the assumption that the uncertainty is the same across all phases of the project. Increases in the

critical cost to completion may be much larger if a project's uncertainty is largely resolved during its early phases.

V. UNCERTAINTY, INSURANCE AND IRREVERSIBILITIES IN CLIMATE POLICY

In the previous sections of this review, I have described a number of Professor Robert Pindyck's most notable contributions. This includes the implementation of optimal control methods to identify optimal economic policies, analyzing how market structure and market power affect the price and output trajectories of non-renewable resources, and the roles of uncertainty, option values, and irreversibilities in affecting investment behavior. In this final review section, I describe how Professor Pindyck has integrated these various concepts and insights into his analysis of economic issues underlying climate policies.

V.A. THE CRITICAL IMPORTANCE OF UNCERTAINTY IN CLIMATE POLICY

Much public attention has focused on the irreversible damages wrought by growing CO₂ emissions and other pollutants that contribute to global warming. Is the existence of growing irreversible damages the principal motivation for enacting climate change policies today? No, writes Pindyck in his most recent book, *Climate Future: Averting and Adapting to Climate Change*¹³; rather, it is uncertainty regarding the possibility of extreme adverse outcomes, and the opportunity to acquire insurance, that leads one to take immediate actions to mitigate climate change.

Regarding consequences of accumulating pollution and failure to take actions to mitigate global warming, Pindyck [2022, p. 75] emphasizes "Rarely do we read or hear that those things *might happen*; instead we're told they *will happen*...The extent of climate change and its impact on the economy and society more generally, is far more uncertain than most people think." What does the existence of uncertainty imply for policies? Pindyck begins his analysis as follows: "You might think so much uncertainty should lead us to wait and see what happens, rather than try to sharply reduce emissions right away. After all, if we don't know how much the climate will change, and we don't know what the impact of climate change will be, why take costly actions now? That is indeed the argument made by many of the people who oppose the imposition of carbon

¹³ See especially his chapter 4, "The Role of Uncertainty in Climate Policy", Pindyck [2022], pp. 75-95.

taxes or other measures to reduce emissions. But that argument is wrong, and actually gets things backward. As we will see, *the uncertainty itself can lead us to act now*. Why? Because with uncertainty, and especially with the possibility of an extreme outcome, we need insurance” (Pindyck [2022, pp. 75-76]).

Irreversibilities are an inherent part of climate policy (and environmental policy more generally). It has long been understood that environmental damage can be irreversible, which can lead to a more “conservative” policy than would be optimal otherwise. And if the value of a cleaner environment to future generations is uncertain, the benefit from protecting it today should include an “option value”, which pushes the cost-benefit calculation towards protection.

A second kind of irreversibility, however, works in the opposite direction: Protecting paradise imposes *sunk costs* on society, costs that cannot be recovered, so that the expenditure is irreversible. To protect clean air and water could require sunk cost investments in abatement equipment, and an ongoing flow of sunk costs for more expensive production processes – funds that cannot be recovered in the future. This kind of irreversibility would lead to policies that are less “conservationist” than they would be otherwise, i.e., they would push the cost-benefit away from protection. While both these irreversibilities are important, we can’t say which one is more important.

Pindyck then argues that while there are uncertainties associated with environmental damages, they are limited and can be bounded. However, when it comes to climate sensitivity, the uncertainty of, say, mild vs. severe temperature outcomes, is much greater. To Pindyck, these uncertainties make the design and analysis of climate policy very different from most other problems in environmental economics, most of which are amenable to standard cost-benefit analysis, where we compare the cost of any particular emission reduction to the resulting benefit, and consider reducing emissions further if the cost is less than the benefit. To Pindyck, “there will be uncertainties over the costs and benefits of any candidate policy, but the characteristics and extent of those uncertainties will usually be well understood and comparable in nature to the uncertainties involved in many other public and private policy or investment decisions...But at a basic level, we’re in well-charted territory and we think we know what we’re doing.” (Pindyck [2022, pp. 77-78]).

But, Pindyck asserts, this is not the case when it comes to climate change. With climate change, there is disagreement among both climate scientists and economists over the likelihood of alternative climate outcomes, especially catastrophic outcomes. There is also disagreement about the framework to employ – for example, the discount rate to be used to compare future benefits with present costs is particularly important in climate policy, because most of the benefits will occur in the far future. These disagreements make climate policy much less amenable to standard cost-benefit analysis. To Pindyck, “The bottom line is that climate policy is complicated by the huge amount of uncertainty we face over the extent and impact of future climate change” which is “especially problematic when it comes to catastrophic outcomes” (Pindyck [2022, p. 78]).

One way to deal with uncertainty is to treat estimated parameters in integrated assessment climate models as random variables, and then observe the sensitivity of model outcomes to various parameter values via probability distributions; this is called *Monte Carlo* simulation. The probability distribution for a given parameter is typically chosen by the modeler, and represents the modeler’s views about the nature of uncertainty for that parameter. With 10 or 20 parameters, each with a probability distribution, the modeler has numerous combinatorial possibilities, and thus might run the model 100,000 times, generating a distribution (with mean and variance) for the output variable of interest, such as lost GDP at the end of the century.

While Monte Carlo simulation can be a powerful tool for incorporating uncertainty in a model, and is widely used, to Pindyck “...it is useful only when applied to a model that has a strong theoretical and empirical foundation, and has parameters for which the probability distributions are well understood and empirically supportable. In the case of climate change, however, we know as little about the correct probability distributions as we do about the damage function to which they are applied. What can we expect to learn from running Monte Carlo simulations? Unfortunately, not much” (Pindyck [2022, p. 81]).

Pindyck is equally skeptical about other methods for incorporating uncertainty into climate models. He quotes Mervyn King, former Governor of the Bank of England, who stated “...if we don’t know what the future might hold, we don’t know, and there is no point pretending otherwise” (Pindyck [2022, fn. 5, p. 81]).

V. B. CLIMATE INSURANCE TO DEAL WITH MASSIVE UNCERTAINTY

What to do given this massive uncertainty? Pindyck steps back and re-examines the value of climate insurance. Modelers of climate change typically utilize a “damage function”. Although the impact of any change in global temperature, for example, is uncertain, Pindyck argues that it is very likely that the damage function becomes increasingly steep; as the temperature change becomes larger, damages become more severe and adaptation becomes more difficult, so the incremental damage from an additional 1° Celsius of warming becomes ever larger (i.e., marginal impacts of temperature change increase with increases in temperature). As a result, for example, a 4° Celsius temperature change causes more than twice the damage of a 2° temperature change. Pindyck then shows how uncertainty, combined with an increasingly steep damage function, creates a value and demand for insurance.

Suppose the percent loss of GDP resulting from a temperature increase ΔT , denoted as $L(\Delta T)$, is equal to the degrees of temperature change squared, i.e.

$$L(\Delta T) = (\Delta T)^2. \quad (\text{Eqn. 21})$$

Eqn. 21 states that $L(0) = 0$, i.e., with no temperature increase, there would be no loss of GDP. It also says $L(2) = 4\%$, i.e., a 2° C temperature increase would result in a loss of 4% of GDP, $L(4) = 16\%$, i.e., a 4° C temperature increase would result in a 16% loss of GDP, $L(6) = 36\%$, i.e. a 6° C temperature increase would result in a 36% loss of GDP, and so on. Note that each additional 2° C increase in temperature results in a larger and larger additional loss, revealing the “increasingly steep damage function”.

Now suppose we knew for certain that in 2050 the global mean temperature will increase by 2° C. And suppose we knew for certain that this 2° C will result in a 4% decline in GDP, compared to what GDP would be without the higher temperature. What percent of GDP should we be willing to sacrifice to avoid this temperature increase? Up to 4 percent, though we’d like to avoid the temperature increase at a cost of less than 4 percent, by developing and making use of energy-saving equipment. But if we needed to, we’d be willing to sacrifice up to a maximum of 4% of GDP.

Now instead of certainty suppose there is uncertainty over the temperature increase – it might not increase at all, or it might increase by 4° C, with each outcome having a 50% probability. The expected value (mean) of the temperature increase equals $0.5 \cdot 0^\circ + 0.5 \cdot 4^\circ = 2^\circ \text{ C}$ – the same as it was in the certainty case, but now there is uncertainty – it might be zero and it might be 4° C.

In terms of its damaging impact on GDP, since a 2° C temperature increase resulted in a 4% decline in GDP, would it be reasonable to assume a 4° C increase would be twice that of a 2° C increase, i.e. 8%? No – because of the increasingly steep damage function discussed earlier, we would expect the damages associated with a 4° C increase in average global temperatures to be much larger – according to the damage loss Eqn. 21, we’d expect it would cause a 16% drop in GDP. In this case, what percent of GDP should we be willing to sacrifice to avoid the possibility of a 4° C temperature increase?

To address that issue, consider the expected size of the impact on GDP. The expected size of the temperature change is still 2° C, and we noted earlier the impact of a 2° C temperature increase would be 4% of GDP. But the expected impact of a fifty-fifty chance of no temperature increase and a 4° C temperature increase is greater than 4% of GDP. Specifically, it is $0.5 \cdot 0\% + 0.5 \cdot 16\% = 8\%$ of GDP. That indicates we should be willing to sacrifice up to 8% of GDP to avoid the 50% chance of a 4° C temperature increase. At 8%, we’d be willing to sacrifice more than 4% of GDP because the 4° C increase in temperature, which admittedly has only a 50% chance of occurring, would be so much more damaging.

Going one step further, suppose instead there is a 75% probability that there will be no temperature increase, and just a 25% chance of an 8° C temperature increase. According to the damage loss Eqn. 21, an 8° C temperature increase would be arguably catastrophic, resulting in a 64% decline in GDP. Note the expected value of the temperature increase is still 2° C (since $0.75 \cdot 0^\circ + 0.25 \cdot 8^\circ = 2^\circ \text{ C}$), but the expected value of this temperature gamble is now much greater than 4% of GDP. By Eqn. 21 above, it is $0.75 \cdot 0\% + 0.25 \cdot 64\% = 16\%$ of GDP.¹⁴

¹⁴ The “increasingly steep damage function” reflects a mathematical phenomenon known as Jensen’s Inequality: In mathematical terms, for any convex function g , $E[g(X)] \geq g(E[X])$, where E is the expected value operator and X is the random variable.

That suggests that if we had to, we'd be willing to sacrifice up to 16% of GDP to avoid a 25% chance of an 8° temperature increase. So we would be willing to pay a lot to avoid a very bad outcome, even if that outcome had only a small chance of occurring.

As Pindyck [2022, pp. 85-86] notes, "This is the essence of what insurance is all about: We pay to avoid a very bad outcome, even if that outcome is unlikely. That is why we insure our homes against major damage from fire, storms, or floods, why we buy medical insurance to cover the cost of major hospitalizations, and why we buy life insurance even if we are healthy and expect to live many more years. And that is why we as a society should be willing to pay a considerable amount for insurance against a very bad (even if unlikely) climate outcome."

V.C. RISK AVERSION AS ANOTHER RATIONALE FOR BUYING CLIMATE INSURANCE

While the presence of an increasingly steep damage function provides a rationale for insuring against global climate warming, another phenomenon – *the declining marginal utility of income* provides a reenforcing motivation. The extra satisfaction we obtain from being gifted \$50,000 is considerably smaller if our initial income is \$1 million than if it is \$25,000. Similarly, a decline in our income of \$50,000 is more harmful to us if our initial income is \$75,000 than if it is \$1 million. And a decline in income of \$50,000 is more than twice as harmful to us as a decline in income of \$25,000. These are all examples of a declining marginal utility of income. Pindyck [2022, p. 87] argues that what we call a declining marginal utility of income corresponds to *risk aversion*.

Another example of risk aversion is the following: You would probably refuse a lottery in which you had a 50-50 chance of winning \$100,000 or losing \$100,000. The reason is that for most people the value of winning \$100,000 is less than the lost value of losing \$100,000. How much would you have to be paid to agree to take part in that lottery? Perhaps \$25,000, so that you'd have a 50-50 chance of winning \$125,000 or losing \$100,000. The greater the amount you'd have to be paid, the greater is your risk aversion. How risk averse society as a whole is difficult to determine, since people have diverse attitudes toward risk. Financial

market data tell us that investors in the aggregate seem to have substantial risk aversion, but not everyone is an investor, and averting climate change is not the same as investing in the stock market.

The joint existence of increasingly steep damage functions and risk aversion highlights why uncertainties over climate change are so important, and in particular why society should be willing to sacrifice a substantial amount of GDP to avoid the risk of an extremely bad climate outcome, even if the risk is small. The risk of an extreme outcome – what some call a “black swan” event or “tail risk” – might compel us to adopt a stringent emission abatement policy quickly, rather than waiting to see how bad climate change turns out to be. By reducing emissions now we would be buying insurance, and the value of that insurance could be considerable.

V.D. ENVIRONMENTAL PROTECTION EXPENDITURES AS IRREVERSIBLE SUNK COSTS

Pindyck [2022, pp. 88-90] reminds us, however, that a second type of irreversibility exists that works in the opposite direction: Protecting the environment imposes sunk costs on society. Keeping our air and water clean requires sunk cost investments in abatement equipment, and an ongoing flow of sunk costs for alternative and perhaps more costly production processes. If in the future clean air and water turn out to be less valuable than we currently expect, we will regret the irreversible expenditures that were made, and that could have been spent on other things. This leads us to less “conservationist” policies.

Climate change involves both types of irreversibilities. Because CO₂ can remain in the atmosphere for centuries, and ecosystem destruction from climate change can be permanent, there is clearly an irreversibility argument for taking early and decisive action. But reducing carbon emissions can be quite costly in terms of reduction in GDP, and those costs are largely sunk, implying a rationale for waiting.

We know both these irreversibilities are important, and they work in opposite directions. Which type of irreversibility will dominate depends in part on the nature and extent of the uncertainties involved – how much environmental damage will result from a higher atmospheric carbon concentration, and how reversible is such damage? Pindyck [2022, p. 89] argues that “The problem is that

we don't know how those future damages will be valued, and we won't know until the damages actually occur in the future". If the damages are irreversible, there is almost no limit to the regret we might feel from not taking action today to reduce emissions and limit those future damages. As a result, argues Pindyck, the benefit from reducing emission should include an "option value", which pushes the cost-benefit calculation toward early action.

Regarding the second irreversibility, the value of the damages might turn out to be only moderate, or slight, or even zero, in which case we would regret having spent resources today to reduce the damages. In that case waiting has an "option value", pushing the cost-benefit calculation away from early action.

V.E. INSURANCE VALUE OF CLIMATE ACTION DOMINATES SUNK COST IRREVERSIBILITIES

Pindyck [2022, p. 90] addresses the question of which of these irreversibilities is more important as follows: "The uncertainties over the effects of emission reductions on temperature change and the effects of temperature change on GDP and welfare are so large that we can't determine the net effect of the two opposing irreversibilities. On the other hand, these very large uncertainties imply that the insurance value of early action is large. Whatever the effects of the irreversibilities, they are likely to be swamped by this insurance value – which pushes us to early action".

VI. CONCLUDING REMARKS

Robert Pindyck's contributions to various economic literatures are remarkably numerous, as I have documented in the previous pages. Although uniformly rigorous, based on strong theoretical foundations and often embedded in dynamic contexts, these contributions are not just theoretical, but also inform and help us understand policy controversies. Pindyck's contributions are much more than just important articles in economics journals; he is the author of two very successful textbooks – one in microeconomics and the other in econometrics; and he has been awarded numerous teaching awards by students at MIT. Robert Pindyck is a consummate educator.

VII. CITED REFERENCES

Caballero, Ricardo J. and Robert S. Pindyck [1996], "Uncertainty, Investment, and Industry Evolution", *International Economic Review*, August, 37(3): 641-662.

Chow, Gregory C. [1972a], "On the Optimal Control of Linear Economic Systems with Finite Time Horizon", *International Economic Review*, February.

Chow, Gregory C. [1972b], "How Much Could Be Gained by Optimal Stochastic Policies?", *Annals of Economic and Social Measurement*, October, 1:391-406.

Chow, Gregory C. [1973], "Effect of Uncertainty on Optimal Control Policies", *International Economic Review*, 14:632-645.

Eisner, Mark and Robert S. Pindyck [1973], "A Generalized Approach to Estimation as Implemented in the Troll/I System", *Annals of Economic and Social Measurement*, 2(1):29-61.

Goldberger, Arthur S. [1964], *Econometric Theory*, New York: John Wiley & Sons, Inc.

He, Hua and Robert S. Pindyck [1992], "Investments in Flexible Production Capacity", *Journal of Economic Dynamics & Control*, 16:575-599.

Hnyilicza, Esteban and Robert S. Pindyck [1976], "Pricing Policies for a Two-Part Exhaustible Resource Cartel: The Case of OPEC", *European Economic Review* 8:139-154.

Hotelling, Harold [1931], "The Economics of Exhaustible Resources", *Journal of Political Economy*, April, 39(2):137-175.

Levhari, David and Robert S. Pindyck [1981], "The Pricing of Durable Exhaustible Resources", *Quarterly Journal of Economics*, August, 96(3):365-377.

Majd, Saman and Robert S. Pindyck [1987], "Time to Build, Option Value, and Investment Decisions", *Journal of Financial Economics*, 18:7-27.

Martens, Andre, and Robert S. Pindyck [1975], "An Optimal Control Model for Multisectoral Investment Planning in Tunisia", *Journal of Development Economics*, 2:99-119.

Nash, John F. [1953], "Two-person cooperative games", *Econometrica*, January, 21:128-140.

Neck, Reinhard [2008], "The Contribution of Control Theory to the Analysis of Economic Policy", Proceedings of the 17th World Congress, The International Federation of Automatic Control, Seoul, Korea, July 6-11. Available online from reinhard.neck@uni-klu.ac.at.

Neese, John Wesley and Robert S. Pindyck [1984], "Behavioural Assumptions in Decentralised Stabilisation Policies", ch. 12 in Andrew J. Hughes-Hallett, ed., *Applied Decision Analysis and Economic Behaviour*, Dordrecht, Netherlands: Martinus Nijhoff Publishers, pp. 251-270.

Pindyck, Robert S. [1970], "A Small Quarterly Model of the U.S. Economy", unpublished manuscript, Cambridge, MA: MIT Economics Department, April.

Pindyck, Robert S. [1971], "Optimal Economic Stabilization Policy", Ph.D. dissertation, Massachusetts Institute of Technology, Department of Economics, June.

Pindyck, Robert S. [1972a], "An Application of the Linear Quadratic Tracking Problem to Economic Stabilization Policy", *IEEE Transactions on Automatic Control*, Vol. AC-17(3):287-300, June.

Pindyck, Robert S. [1972b], "Optimal Stabilization Policies via Deterministic Control", *Annals of Economic and Social Measurement*, 1(4):385-389.

Pindyck, Robert S. [1973a], *Optimal Planning for Economic Stabilization: The Application of Control Theory to Stabilization Policy*, Amsterdam: North-Holland Publishing Co.

Pindyck, Robert S., [1973b], "Optimal Policies for Economic Stabilization", *Econometrica* May, 41(3):529-560.

Pindyck, Robert S. [1976], "The Cost of Conflicting Objectives in Policy Formulation", *Annals of Economic and Social Measurement* 5(2):239-248.

Pindyck, Robert S. [1977], "Optimal Economic Stabilization Policies Under Decentralized Control and Conflicting Objectives", *IEEE Transactions on Automatic Control*, August, 22(4):517-530.

Pindyck, Robert S. [1978a], "Gains to Producers from the Cartelization of Exhaustible Resources", *Review of Economics and Statistics*, May, 60(2): 238-251.

Pindyck, Robert S. [1978b], "The Optimal Exploration and Production of Nonrenewable Resources", *Journal of Political Economy*, 86(5):841-861.

Pindyck, Robert S. [1980], "Uncertainty and Exhaustible Resource Markets", *Journal of Political Economy*, December, 88(6):1203-1225.

Pindyck, Robert S. [1984], "Uncertainty in the Theory of Renewable Resource Markets", *Review of Economic Studies*, 51:289-303.

Pindyck, Robert S. [1985], "The Measurement of Monopoly Power in Dynamic Markets", *Law & Economics*, April, 28(1):193-222.

Pindyck, Robert S. [1986], "Capital Adjustment Costs, Monopoly Power, and the Regulated Firm", chapter 11 in R. E. Grieson, editor, *Antitrust and Regulation*, Lexington, MA: Lexington Books, 219-230.

Pindyck, Robert S. [1988], "Irreversible Investment, Capacity Choice, and the Value of the Firm", *American Economic Review*, December, 78(5):969-985.

Pindyck, Robert S. [1993a], "A Note on Competitive Investment under Uncertainty", *American Economic Review*, March, 83(1):273-277.

Pindyck, Robert S. [1993b], "Investments of Uncertain Cost", *Journal of Financial Economics*, 34:53-76.

Pindyck, Robert S. [1994], "Inventories and the Short-Run Dynamics of Commodity Prices", *RAND Journal of Economics*, Spring, 25(1):141-159.

Pindyck, Robert S. [2022], *Climate Future: Averting and Adapting to Climate Change*, New York, NY: Oxford University Press, 230 pp.

Pindyck, Robert S. and Steven M. Roberts [1974], "Optimal Policies for Monetary Control", *Annals of Economic and Social Measurement*, 3(1):207-236.

Pindyck, Robert S. and Steven M. Roberts [1976], "Instruments, Intermediate Targets, and Monetary Controllability", *International Economic Review*, October, 17(3):627-650.

Pindyck, Robert S. and Julio J. Rotemberg [1983a], “Dynamic Factor Demands under Rational Expectations”, *Scandinavian Journal of Economics*, 85(2):223-238.

Pindyck, Robert S. and Julio J. Rotemberg [1983b], “Dynamic Factor Demands and the Effects of Energy Price Shocks”, *American Economic Review*, December, 73(5):1066-1079.

Pindyck, Robert S. and Daniel L. Rubinfeld [1998], *Econometric Models and Economic Forecasts*, Fourth Edition, Boston: Irwin McGraw-Hill.

Pindyck, Robert S. and Daniel L. Rubinfeld [2013], *Microeconomics*, 8th edition, Boston, Pearson Education Inc., publishing as Prentice Hall.

Pontryagin, Lev Semenovich, Vladimir Grigorevich Boltyanskii, Revaz Valerianovic Gamrelidze, and Evgenii Frolorovich Mishchenko [1962], *The Mathematical Theory of Optimal Processes*, editor L. W. Neustadt, translated from Russian to English by Trirogoff. Interscience Publishers, New York-London, John Wiley & Sons, inc. MR 0166037 (<https://www.ams.org/mathscinet-getitem?mr=0166037>).

“Professor Emeritus Michael Athans, Pioneer in Control Theory” [2020], Obituary, June 8. Available online at <https://news.mit.edu> > professor-emeritus-michael-athan... Last accessed 9 February 2023.

Theil, Henri [1957], “A Note on Certainty Equivalence in Dynamic Planning”, April, *Econometrica*, 25:346-349.

DOCUMENT NAME: PINDYCK FULL DRAFT

DOCUMENT DATE: 22 MARCH 2023