The Rate, in Its Place

Policymakers facing environmental policy with a long horizon should apply a low, positive, declining social discount rate: around 2% in the near term, drifting toward 1% or under on a century scale. That figure is roughly the median of recent expert opinion (Drupp, Freeman, Groom and Nesje 2017), and the schedules already in operational use in France and the UK look much the same. So far, the simple part. The harder claim is that the social discount rate has a domain of competence. The rate compares; it does not unproblematically choose. Four assumptions sit underneath it. The project is a marginal perturbation of an otherwise stable consumption path. The goods traded across time are substitutable enough that a single welfare numeraire is meaningful. The relative prices used to value impacts are stable, rather than themselves the object the policy is meant to alter. And the impact distribution is well-behaved enough for ordinary expected-value methods to apply. Some environmental policies break those assumptions. They are non-marginal; they foreclose choices that cannot later be re-opened. A tipping-point lock-in of the climate system, or the exhaustion of a non-substitutable natural resource: once such a policy is the object of analysis, the rate becomes endogenous to it. The right response is not to keep tuning the number; it is to change the instrument. Impose an exogenous admissibility constraint, and discounting can then do what it was built to do, which is to compare policies that satisfy it. The rest of this essay derives the rate and traces the boundary that defines its proper use.

The standard formula is the Ramsey rule, $r = \delta + \eta g$: the social discount rate r equals pure time preference δ plus the elasticity of marginal utility of consumption η, multiplied by expected per-capita consumption growth g (Ramsey 1928). The canonical split sets the Stern Review against Nordhaus. Stern (2007) sets δ near zero on broadly impartial-utilitarian grounds and arrives at a rate near 1.4%. Nordhaus (2007), reading the rate off observed market returns and savings behaviour, lands at 4–5%. Each position survives in the face of the other. Stern's rate looks too low when set beside Nordhaus's emphasis on capital productivity. Nordhaus's looks too high when set beside the moral revulsion at discounting a 1985-born person at half the value of a 1960-born one. To treat the debate as a fight over Ramsey parameters, though, mistakes the question. The more useful anchor is empirical. An expert survey put the median recommended social discount rate at 2%, with most responses clustered between 1% and 3% (Drupp, Freeman, Groom and Nesje 2017). That is the figure to start from.

Two per cent should not, however, be held fixed across the horizons that matter. The certainty-equivalent term structure of interest rates bends downward over long periods

whenever there is both persistent uncertainty about long-run growth and aversion to long-run risk (Gollier 2012). The reasoning is direct. Persistent shocks to the growth rate make the variance of consumption at time t a rising function of t. That sharpens the precautionary motive for saving. The correct rate for a far-future payment then sits below the correct rate for a near one. Weitzman's (2001) gamma-discounting result, drawn from a divergence in expert views, points the same way. The case for a declining schedule on long-horizon public policy is made by Cropper, Freeman, Groom and Pizer (2014) and by the expert panel of Arrow et al. (2014); France and the UK already have one in operational use. So the rate declines. That part is settled.

The next proposition is more contested. A social discount rate is a tool for making comparisons; it turns intertemporal streams of value into a single present-value figure. Two assumptions sit beneath that machinery, and both are usually left implicit. The first is marginality: the project shifts the consumption path only slightly, so the relative prices used to value its impact can be read off the unperturbed path ex ante. The second is substitutability: the goods traded across time are interchangeable enough that the trade-offs between them collapse into a common welfare numeraire. Both hold for the bulk of long-horizon problems, including the standard examples: sea-wall provision, the harvesting of timber, a carbon tax imposed on the baseline scenario. What follows is a triage rule, not a replacement for the three structural failures set out below.

The boundary admits a simple test. A long-run environmental policy falls outside the standard discount-rate framework when at least two of four conditions hold. The policy is non-marginal. Its impact is effectively irreversible. The resource at stake is poorly substitutable. The impact distribution is too fat-tailed or ambiguous for ordinary expected-value methods to capture. Each flag is well known from the discounting literature on its own; what the diagnosis adds is the combination. The four are not a prescription, only a triage. They tell us whether a rate of discount is the right basis for admissibility, or only the right basis for choosing among alternatives once admissibility has already been settled.

The top row is the rate's home ground. The bottom row is not. The middle row is where careful work earns its keep. The point is structural rather than numerical: the ordinary case is a choice within the paradigm of discounting, and the boundary cases need sensitivity and robustness analysis on top of that. But the issues that drive political salience around Stern–Nordhaus sit squarely in the regime that needs a constraint imposed first. The decades-long argument over what number to attach to r was not so much wrong as misdirected.

Three distinct and interacting problems sustain the boundary. First, relative prices. If consumption goods and environmental services are not perfect substitutes, what is needed is not a rate but a term structure of relative prices. The relevant price path is then conditioned on the scarcity of the environment (Sterner and Persson 2008; Gollier

Policy decision	**Warning signs triggered**	**Diagnostic result**
Municipal bus electrification	None or one: marginal, mostly reversible, substitutable	Discounting decides
Wetland restoration at ecosystem scale	One or two depending on scale: partial irreversibility, limited substitution	Borderline / constrained CBA
1.5–2°C climate pathway choice	Non-marginality, irreversibility, limited substitution, deep uncertainty	Constraint first

2012, ch. 10), and the path is itself partly a function of the foreclosure decision under examination. The single-rate analysis becomes conditioned on assumptions about the very price path the decision is meant to alter. No value of η can carry that load. What is missing is not a number, but a path. Second, the framework's own internal admission of the difficulty. Gollier (2012, ch. 14) argues in *Pricing the Future* that marginalist pricing fails for non-marginal projects: such projects shift the consumption path and the equilibrium prices along it, so the relevant discount rate and risk premium become endogenous outputs rather than exogenous inputs. He adds that simply lowering g in the Ramsey rule to reflect environmental damages will not do, because the policy under examination is meant to change both the intensity and the variance of climate change (Gollier 2012, ch. 14). A rule built for pricing small adjustments along a given path cannot price a decision whose purpose is to alter that path. Third, the difficulty is structural rather than numerical. Discounting compares options inside a feasible set; irreversible foreclosure works by removing options from that set. The rate cannot value an option it has just deleted. Dasgupta (2008), looking at intergenerational welfare under climate change, has shown that under relatively plausible assumptions an optimum may simply fail to exist: the problem is not just badly behaved but mathematically ill-posed. The point is not an external moralistic objection; it is a shortcoming flagged from inside the

framework, by an economist friendly to it.

The strongest reply from inside the paradigm is that the problem can be calibrated out. The severe effects of environmental loss are, the argument goes, absorbable into the standard kit: a low discount rate, a high η asked to carry inequality aversion, risk aversion, and intertemporal substitution all at once, with a damage function convex enough to take care of the rest. All of this is available, and within the marginal regime calibration does real work. That part I do not contest. What I do contest is the move from there to the boundary cases. Calibration adjusts numbers; it cannot fix a mismatch between the tool and the decision it is applied to. Where policy is itself rewriting the relative-price trajectory, the object is no longer a static figure but an evolving path, and the path is endogenous by construction. Where a project is non-marginal in Gollier's sense, the marginal pricing rule does not apply, and tuning its parameters changes nothing about the rule. Where an option has been deleted from the feasible set, the rate has no machinery for the opportunity cost it imposes; that absence is formally registered by Dasgupta's non-existence theorem. To calibrate harder is, in effect, what Nordhaus accuses Stern of doing. Both poles miss the point. The diagnosis is structural, and structural mismatches do not yield to tuning.

A harder pragmatic objection cuts the other way. Constraints carry trade-offs too. To refuse to price a good can become a price by another name, only higher, less visible, and less disciplined than a transparent social cost of carbon. The objection is fair as far as it goes. A binding admissibility limit is, after all, itself a trade-off mechanism. The question is where the trade-off lives. A single present-value figure hides the real trade-off inside η, inside the damage function, and inside an assumed substitutability of environmental goods that is doing most of the work without being asked to defend itself. A binding constraint puts the premise back on the surface (these emissions, this temperature, this minimum stock are out of bounds), and the discount rate then locates the least-cost admissible path within those terms. Nordhaus is right about opportunity cost, productivity, and abatement comparisons. He is right after admissibility has been settled, not before. The trade-offs themselves do not vanish under a constraint; they move. They live in the choice that fixes the price path.

The response is therefore two-track. For ordinary environmental project appraisal, governments should use a low, positive, declining social discount rate, around 2% in the near term and falling toward 1% or under at the century scale: consistent with where the expert distribution sits, and consistent with what we know about the uncertainty profile of long-run growth. For projects that are non-marginal in Gollier's sense (i.e., they shift the trajectory toward a different optimum rather than tweaking it), that are at least partly irreversible on policy-relevant timescales, and whose value turns mainly on the price trajectory of currently unpriced environmental services, admissibility should be set by a binding constraint, with the discount rate used only to choose among feasible

implementations of it. The test is narrow rather than mechanical, but it tracks practice: most environmental projects sit comfortably inside the discounting domain, and the small remainder is where careful policy already concentrates its attention. Discounting can compare what is on offer. Only a constraint can keep an option from being deleted before the rate ever sees it. Whether the future retains its option set is not a question a price of time can settle; the rate is allowed to choose only among the options a constraint has already let through.

References

Arrow, K. J., Cropper, M. L., Gollier, C., Groom, B., Heal, G. M., Newell, R. G., Nordhaus, W. D., Pindyck, R. S., Pizer, W. A., Portney, P. R., Sterner, T., Tol, R. S. J., and Weitzman, M. L. (2014) "Should governments use a declining discount rate in project analysis?" *Review of Environmental Economics and Policy*, 8(2), 145–163.

Cropper, M. L., Freeman, M. C., Groom, B., and Pizer, W. A. (2014) "Declining discount rates," *American Economic Review: Papers and Proceedings*, 104(5), 538–543.

Dasgupta, P. (2008) "Discounting climate change," *Journal of Risk and Uncertainty*, 37, 141–169.

Drupp, M. A., Freeman, M. C., Groom, B., and Nesje, F. (2017) "Discounting disentangled," Memorandum 20/2015, University of Oslo, Department of Economics.

Gollier, C. (2012) *Pricing the Future: The Economics of Discounting in an Uncertain World*. Princeton, NJ: Princeton University Press.

Nordhaus, W. D. (2007) "A review of the Stern Review on the economics of climate change," *Journal of Economic Literature*, 45(3), 686–702.

Ramsey, F. P. (1928) "A mathematical theory of saving," *Economic Journal*, 38, 543–559.

Stern, N. (2007) *The Economics of Climate Change: The Stern Review*. Cambridge: Cambridge University Press.

Sterner, T. and Persson, U. M. (2008) "An even Sterner review: introducing relative prices into the discounting debate," *Review of Environmental Economics and Policy*, 2(1), 61–76.

Weitzman, M. L. (2001) "Gamma discounting," *American Economic Review*, 91(1), 260–271.