

in facts more than in values, producing relatively objective assessments of what would and what would not work for a given problem.⁴

Many aspects of this vision are firmly entrenched as mainstream policy analysis. Graduate courses in policy analysis typically train students to view policy analysis as a linear problem-solving paradigm whose main objective is to advise a client on the best policy alternative for a given problem (Shulock 1997, 228; Weimer and Vining 2005, 1; Bardach 2005/2011; Patton, Sawicki, and Clark 2013, 3–6). This is often accomplished using sophisticated quantitative methods; just about everything in the statisticians', econometricians', and game theorists' tool kit has been adapted to serve the needs of policy analysis (for a comprehensive introduction to the range of quantitative policy analysis methods, see Gupta 2001/2010; Dunn 1981/2011). The increasing sophistication of methods, however, has not yet addressed a fundamental conceptual problem that the rationalist project has struggled with since its modern inception.

If the underlying purpose of the rationalist project is to come up with some ranking of policy alternatives, it needs a measure with which to rank them. On what grounds does rational policy analysis judge a particular policy alternative "best"? In Franklin's system of moral algebra this was a comparatively easy question to answer. Because the problems were individual problems, the subjective values of the individual were enough to attach weights to the pros and cons of various choices. The rationalist policy analyst does not have it so easy; the subjective values of the analyst are not a valid yardstick to judge the worth of public policies that represent collective interests. But what values should the analyst use to determine which policy alternative is the best response to a given problem? Answering this question is a central challenge for the rationalist project.

The generic problem-solving process listed above requires a clear definition of the problem to be addressed and accurate forecasts of how differing policy options will address that problem. In practice, however, the problems that public policies are called upon to solve rarely have a universally agreed-upon, precisely defined, easily measurable, single goal. Educational performance, for example, is a pretty vague goal. How is it measured? With test scores? Graduation rates? If that question is answered comprehensively (a tall order in and of itself), how do we estimate the impact of different policy options on test scores or graduation rates? We cannot observe what has not happened, so we must somehow accurately forecast the impact of policy alternatives on these measures. And even if we manage to do that (an even taller order), a policy analysis is still likely to get pretty messy. For example, it may show that the program forecast to increase test scores the most—the one with the tougher academic standards, say—lowers graduation rates, and that the policy that increases graduation rates the most is forecast to lower test scores. Which option should we choose? What if the policy