

Optimal Confidence Procedures for Partially Identified Parameters

Ian Pitman

October 23, 2026

Abstract

This paper constructs optimal confidence procedures in a flexible class of partially identified models. First, I develop a decision-theoretic framework in which reporting a confidence set incurs a loss equal to its excess length: the length of the part that does not overlap with the identified set. Under my proposed framework, a confidence procedure is optimal if it minimizes worst-case expected excess length. Next, I show that the problem of constructing such a procedure reduces to solving a single optimal testing problem and inverting the resulting test. This theoretical insight allows me to leverage recent computational advances in optimal testing to numerically approximate the optimal procedure. In a series of calibrated simulations, I quantify the suboptimality of existing inference methods across a variety of empirical settings. I find that methods with the worst performance often direct power towards alternatives that are inconsistent with the model, which is never optimal. Finally, I illustrate the usefulness of my approach in an event study design from Iovino et al. (2026).