Supplemental Appendix for "Analyzing Computational Models"

David A Siegel

Sequential Parameter Sweeping More Formally

I provided an informal description of sequential parameter sweeping (SPS) in the text, making use of two examples to illustrate the method. In this appendix I formalize, to the extent possible, some of its assumptions and arguments. This formalization is useful in applying SPS to more complex models.

I begin by defining one's ideal model, M. This model contains all complexity the modeler believes is necessary to properly instantiate her theory. Since our focus is on the comparative statics of the steady states of computational models, I formally define M as a mapping taking a set of parameters, B, to a set of steady state variables, X. Each individual parameter is β_k, and each individual steady state variable is x_l.

Next I define a sequence of nested models indexed by i, M_i, with $M_m = M$ for some finite m. These have a very specific property: each M_i is nested within all M_j with $j > i$. Formally, to be nested means, for all $j > i$: (i) the set of parameters, B_i, in each model M_i is a strict subset of the set of parameters, B_j, in each M_j; and (ii) the set of steady-state variables, X_i, in each model M_i is a (not necessarily strict) subset of the set of steady state variables, X_j, in each M_j. Practically, this means that one forms nested models by removing one or more parameters and possibly one or more steady-state variables from a more complex model. This nesting is not typically unique, and should be informed by the substantive question. In the interpersonal trade example in the text, the PD addressed a substantively distinct question on its own, and so made sense as a nested component of the more complex model. In the program evaluation example in the text, collective action absent government is one question, the effect of government interventions on collective action is a second, and the effect of bureaucracy on the effect of interventions is a third.

Finally, I define a separability property that connects the analyses of the M_i. Let B_{in}, $n = 1 \ldots N_i$, N_i finite for all i, be disjoint regions of the parameter space of model M_i, such that the union of the B_{in} is the parameter space spanned by the set of parameters B_i. In other words, the B_{in} partition the parameter space of a model M_i. We say the parameter space of model M_i is *separable into regions of common behavior* if three conditions are met.

The first condition is that within each region B_{in} the functional dependence of each of the steady state variables in X_i on each of the parameters in B_i must "look the same." Here "look the same" refers only to the *shape* of the comparative statics within each region, and not to the boundaries between regions or the precise values of each endogenous variable within regions. That is to say, what is important is that the behavior of the model M_i within each region, as represented by the manner in which the endogenous variables change with the exogenous parameters, is similar within the region. One must care about both sign and rate of change of the comparative statics in assessing this condition, and extant theory is particularly helpful in justifying adherence to it. The two examples in the text illustrate what it means to satisfy this condition.

The first condition is, to some extent, deliberately vague: the intent is to capture different regions of the parameter space within which the causal mechanisms driving the dependence of each steady state variable on its parameters are common to the region. As this deliberate vagueness comes up at several points in what follows, it is worth briefly elaborating on its intent before moving on to the second condition.

The intent of the first condition, and SPS in general, is to understand patterns of behavior, rather than point predictions. This same intent underlies comparative statics of game-theoretic models, and SPS is similar to the spirit of one method of deriving them, monotone comparative statics.[1]

To pursue this intent, we must define a standard for what constitutes a pattern of behavior. One is typically only concerned with the signs of comparative statics in game theory models. A pattern of behavior according to this standard is a specification, possibly condi-

[1] Thanks to Ken Shotts for the analogy.

tional, of the signs of all comparative statics across all subsets of the full parameter space.

SPS uses a somewhat different standard: it requires that not only the sign, but also the rough shape of the comparative statics must be similar within each region of common behavior. This standard is stricter than that typically used in game theory in the sense that it is not just the sign that matters, but it is weaker in that the best one can do in a computational model, particularly when models exhibit substantial stochasticity, is to specify a numerical tolerance for what constitutes rough similarity.

The phrase"look the same" arises from this consideration of numerical tolerance. I could have replaced this with a more formal definition along the lines of "the integrated difference of the comparative statics taken at two different parameter values within a region of common behavior must not exceed y across the domain of the region" or "the sign of the comparative statics should not cross more than y beyond zero from above to be considered positive." However, this would have reduced flexibility too much for the end-user of the method, who should remain free to specify whatever standard of "looking the same" is justified by the research context in question.

An end-user of the method could vary the standard of numerical tolerance embedded in SPS in more substantive ways if desired. In some cases, the analyst may find only the sign of the comparative statics of interest, and possess broad tolerance as to numerical variability near zero. In others, the analyst may want a more precise depiction of the complete functional form of the comparative statics, with relatively little error. The former might occur when one's model is intended to be suggestive of an underlying behavior, while the latter might occur when one believes one's model has substantial ground truth, perhaps because many parameters have been fitted to data. Somewhere between the two is the standard I have used in this presentation, which was chosen under the presumption that the shapes of the comparative statics capture the most important patterns of behavior to the researcher.

However, when considering changing the standards of SPS more substantially, it is important to note that the standard used does affect the degree to which one can employ SPS successfully. A stricter standard of "looking the same" will generally imply less ability to discover regions of common behavior, and SPS will generally fail more quickly under this stricter standard. In contrast, a very weak standard that merely focuses on the sign of the comparative statics rather than their shape would fail less often than the one I have presented.

I now return to discussing the separability property. To specify its second condition, we need to link nested models. Recall that the shapes of the comparative statics, rather than the precise parameter boundaries that specify each B_{in}, are what truly define each region. Therefore, define regions B'_{jn} of the parameter space of model M_j, $j > i$, in the following manner: In each B'_{jn}, the comparative statics for each parameter in B_i, for every steady state variable in X_i, share roughly the same functional forms as they do in the corresponding region B_{in} of the parameter space of the simpler model M_i. In other words, the comparative statics whose similar behavior defined the B_{in} in model M_i also define the B'_{jn} in model M_j for any $j > i$. Note that because the more complex model M_j has a higher-dimensional parameter space than M_i, the B'_{jn} need not be unique, since they are specified only by the similar behavior of comparative statics with respect to parameters and steady state variables present in the simpler model M_i.

The second condition is that regions B'_{jn} exist for each model M_j, for all $j > i$, such that within each region B'_{jn} the functional dependence of each of the steady state variables in X_i on each of the parameters in B_i must "look the same," as in the first condition. Put another way, our first two conditions on separability of a model M_i imply that as we build up more complex models from M_i, we must still be able to find regions of the parameter space for every model in which all comparative statics that were present in M_i continue to betray roughly the same functional dependence within each region as they did in M_i.

The third and final condition is the most stringent, as it specifies a particular kind of independence between models. It states that within each of the B'_{jn} that satisfy the second condition, for each model M_j, for all $j > i$, the functional dependence of each of the steady state variables in X_j on each of the parameters in B_j that are not also in B_i must be approximately independent of the parameters in B_i. As in the first two conditions, this is left deliberately vague, for the same reasons as above. That said, approximately independent admits a relatively easy empirical test: if one's comparative statics for parameters that only exist in model M_j vary substantially with parameters that exist also in M_i for $j > i$, beyond variation due to being in different regions B'_{jn}, then this condition is violated. A cross-partial derivative of an endogenous variable with respect to a parameter in B_i and a parameter that is in B_j but not B_i that is significantly different from zero at a point within the interior of a region B'_{jn} would indicate a violation of this condition.[2] In other words, the comparative statics with respect to the new parameters added in any more complex model cannot depend significantly on the parameters in the simpler model M_i, beyond any dependence captured by presence within a region B'_{jn} of the parameter space of M_j.

Put slightly differently, the third condition specifies that the effects of all parameters that are in more complex models M_j, $j > i$, but not in the simpler model, M_i, are approximately independent of all parameters in M_i, apart from the role of M_i's parameters in determining the boundaries of the regions of the parameter space. This way of phrasing connects separability to labeling. A nested series of models possesses separability if one can find regions of the parameter space of each M_i, $i < m$, that can be labeled; additional parameters in M_j, $j > i$, but not in M_i can have label-dependent effects on steady-state variables, but not significant effects dependent on the particular values of the parameters in M_i beyond this. Identification of these regions can require both computational characterization of M_i and extant theory.

With separability defined, we can move to specifying SPS. SPS begins by instantiating a simplified version of the model, M_0, with no more than 3 parameters. The reduced number of parameters in M_0 should allow one to completely specify comparative statics for each model parameter in all equilibria/steady states/limiting distributions of M_0. This can be done by direct computation. SPS then iterates the three steps below in order to instantiate a particular sequence of models, M_i, that increase in complexity. This iteration continues until

[2]Leaving open what counts as "significantly different" from zero maintains the needed flexibility discussed above.

either: (i) $M_i = M$ for some i, implying the full model is analyzable with SPS; or (ii) step one fails for some i, implying only part of the model is analyzable with SPS. In the latter case one could either: stop at the most complex model analyzable by SPS; analyze the full model using a sampling procedure, supplementing insights from this analysis with those arising from the partial SPS approach; use a non-nested modeling approach that implicitly ignores interactions between model components; or reformulate one's model in a more hierarchical fashion so as to make more effective use of SPS. The steps to be iterated are as follows, with $i = 0$ in the first pass and M_0 assumed fully analyzed.

1. If possible, separate the parameter space of model M_i into regions of common behavior (see definition above).[3]

2. If step one yields no such regions, or an unmanageably large number of them, cease SPS. Otherwise, iterate model M_i toward model M. Call this new, more complex model M_{i+1}; it should possess no more than 1-2 additional parameters. This iteration should be chosen to make step three possible.[4]

3. Fully characterize all possible equilibria/steady states/limiting distributions of this more complex M_{i+1} and compute comparative statics for each of its new variables and parameters (i.e., variables and parameters that are in M_{i+1} but not M_i). Comparative statics must be computed within each of the regions identified in step one, at representative values of the parameters in M_i that distinguish these regions. In other words, for every labeled region from step one, one must choose a set of representative values of the parameters in M_i that delineates that region; these representative parameters

[3]Multiple such partitions may exist. These can be distinguished on the basis of substantive considerations (e.g., which partition has the clearest substantive meaning), or on which allows SPS to proceed the furthest.

[4]Ideally the order of adding complexity would not matter. However, should separability fail before model M is reached, one must decide which aspect of complexity to add first. Again, this decision can be made on the basis of substantive considerations (e.g., which is the most substantively interesting aspect of complexity), or on which order allows SPS to proceed the furthest. This decision can also be influenced by what approach one desires to take with respect to analyzing M. For instance, if one plans to analyze a series of non-nested models due to the inability of SPS to fully analyze M, then the order should be chosen to enable these non-nested models to tell the clearest story.

are then used to compute comparative statics for the new parameters added in step two. If $M_{i+1} = M$ this step completes SPS; if not, let $i = i + 1$ (i.e., increment the index i by 1) and return to step one.