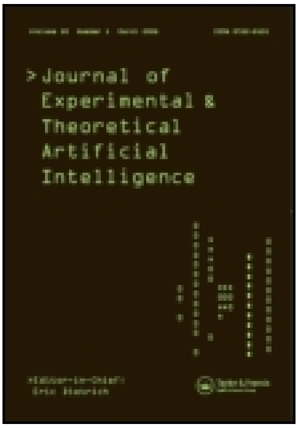

Journal of Experimental & Theoretical Artificial Intelligence
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/teta20

From theory to simulation: the dynamic political hierarchy in country virtualisation models

Ian S. Lustick [a b] , Brandon Alcorn [b] , Miguel Garces [b] & Alicia Ruvinsky [c]

[a] University of Pennsylvania, 3451 Walnut Street, Philadelphia, PA 19104, USA

[b] Lustick Consulting, 10 E. Athens Ave. Suite 212, Ardmore, PA 19003, USA

[c] Lockheed Martin Advanced Technology Laboratories, 4350 Fairfax Dr. #470, Arlington, VA 22203, USA

Version of record first published: 04 Sep 2012.

To cite this article: Ian S. Lustick , Brandon Alcorn , Miguel Garces & Alicia Ruvinsky (2012): From theory to simulation: the dynamic political hierarchy in country virtualisation models, Journal of Experimental & Theoretical Artificial Intelligence, 24:3, 279-299

To link to this article: http://dx.doi.org/10.1080/0952813X.2012.693841

Journal of Experimental & Theoretical Artificial Intelligence
Vol. 24, No. 3, September 2012, 279–299

Taylor & Francis
Taylor & Francis Group

From theory to simulation: the dynamic political hierarchy in country virtualisation models

Ian S. Lustick[ab*], Brandon Alcorn[b], Miguel Garces[b] and Alicia Ruvinsky[c]

[a]*University of Pennsylvania, 3451 Walnut Street, Philadelphia, PA 19104, USA;* [b]*Lustick Consulting, 10 E. Athens Ave. Suite 212, Ardmore, PA 19003, USA;* [c]*Lockheed Martin Advanced Technology Laboratories, 4350 Fairfax Dr. #470, Arlington, VA 22203, USA*

(*Received 26 March 2012; final version received 29 April 2012*)

This article suggests that computer-assisted agent-based modelling has the ability to move beyond abstract representations of political problems to theoretically sound virtualisations of real-world polities capable of producing probabilistic forecasts from distributions of stochastically perturbed model trajectories. In contrast to statistical approaches, this technique encompasses both prediction and explanation, with every distinctive trajectory traceable backwards from the occurrence or non-occurrence of an event of interest through the branching points and mechanisms that led to it. In this article, we illustrate our technique for building a country-scale model from corroborated theories, focusing on the 'dynamic political hierarchy' module that integrates theories of cross-cutting cleavages, nested institutions and dynamic loyalties. We present our forecasts for significant political events in Thailand for the year August 2010–July 2011. Drawing on this case we demonstrate how the challenges of internal validity can be met in complex formal models and conclude by emphasising the importance of advances in visualisation techniques for parsing large amounts of interrelated time-series data.

Keywords: simulation; ABM, scenario analysis, Thailand, virtualisation, forecasting, modelling

1. Introduction

In policy-making circles, the production of multiple 'scenarios' has been a standard technique of analysis for decades. Popularised by Hermann Kahn and the Research and Development (RAND) Corporation as exercises to help think about the 'unthinkable', the practice was widely adopted in industry after Shell Oil successfully used scenario building instead of extrapolation of past trends to anticipate the consequences of the kind of steep rise in oil prices associated with the 1973 oil embargo.[1] The standard technique is to assemble a panel of experts or decision makers and then require each to develop a story about the future focused on 'critical uncertainties' and the 'drivers' (most impactful causal agents) each believes to be most crucially operative. These 'scripts' are then treated as possible futures that might unfold, in part or in whole, as the actual future.

Despite its popularity, however, the approach has significant drawbacks and even dangers. Typically, there are few checks imposed on the imagination of participants, and

*Corresponding author. Email: ilustick@gmail.com

ISSN 0952–813X print/ISSN 1362–3079 online

http://dx.doi.org/10.1080/0952813X.2012.693841
http://www.tandfonline.com

no scrutiny of whether the narratives produced are in fact consistent with the antecedents that have been stipulated as the givens of the exercise, with the conditions that would have been necessary to bring into existence prominent as-yet-unrealised elements in proposed scenarios, or with the details of other scenarios generated by panelists and treated as equally possible by the participants.

As Philip Tetlock (2005) has argued in his celebrated study, *Expert Political Judgment*, scenario building is dangerously vulnerable to a range of cognitive and other biases.[2] One of these biases is the natural tendency to find persuasive stories that contain more detail. This is the problem made famous by Nissim Taleb in *The Black Swan* (2007). However unsystematically generated or weakly substantiated as a forecast a scenario may be, arbitrary and vastly improbable details embedded within in it will powerfully increase its persuasiveness. This is one way in which scenario building ignores the daunting problems associated with disciplined counterfactual thinking. It produces stories about the future that place a premium on rhetoric for their persuasiveness. Typically, these dramatic and interesting stories shift group focus towards a tiny sample of possible worlds based not only on the particular political, data-access, or cognitive biases of individual scenario builders but also on the seductive use of substantively irrelevant, granularly unknowable, but fact-like details (Taleb 2007).[3]

In a series of studies Tetlock and his collaborators have investigated just how well experts in world affairs do if asked to make forecasts in their areas of expertise (Tetlock 1999; Tetlock and Lebow 2001; Parker and Tetlock 2006; Tetlock and Parker 2006). Their results were somewhat dispiriting in at least two ways. First, expert prediction, on the whole, was shown to perform little better than dart-throwing. Second, when confronted with evidence of erroneous forecasts, experts were strongly inclined to find *ad hoc* justifications for not amending their beliefs or theories based on the errors they had made. On the other hand, in *Expert Political Judgment*, Tetlock (2005) was able to show that experts with styles of thinking emphasising expectations of complexity, differentiation and indeterminacy ('foxes'), performed better than doctrinaire theorists ('hedgehogs') and better than dart-throwing. Supporting established findings, he also reported that an average of judgements by a large number of experts outperforms any one expert engaged in a series of forecasts. Thus, even if Tetlock's main conclusion points to the need for some systematic basis for judging the rumination-based forecasts of experts, his work also suggests that expert knowledge is real and that it has the potential to be tapped more effectively for making correct forecasts in human affairs, and specifically world political affairs.

This article offers a brief account of how computer-assisted agent-based models (ABMs) have been used to accomplish this task. The overall strategy is to use the imperatives and capabilities of computer models built to discipline unregulated 'BOGSAT'[4] techniques. Contrary to structural equation or standard game theoretic models, computer-assisted ABMs can accommodate complexity, large numbers of actors and dynamic bottom-up processes. These features enable these models to exploit corroborated (but complex) theories of the social world via virtualised dynamic renderings of political arenas incorporating best available theory, expertise and data. The objective is to replace rhetorical seduction or pre-analytic biases as the criteria for evaluating forecasts with controlled and replicable studies of patterns in large distributions of equally possible, but unequally plausible futures.

The discipline associated with this strategy for producing forecasts of 'events of interest' in particular countries is in part directly a function of the imperatives of computer use. A computer program runs without human discretion. Computer models are therefore completely intolerant of ambiguity and establish the cotenability of all propositions within operationalised theories incorporated into the model. If subjected to random perturbations or Monte Carlo techniques of parameter variation, repeated trials will produce futures or forecasts that may differ, but which will all, automatically, be known as consistent with the theoretical positions instantiated in model design and with initial data inputs. By varying initial data inputs or variables of interest, 'what-if' experiments involving comparison of batches of 'futures' can be conducted. By comparing batches of 'futures' produced under controlled conditions, the robustness of forecasts across a range of possible assumptions can be systematically evaluated.

This kind of research, entailing the virtualisation of a particular problem so that it can be investigated in a controlled manner, is an ambitious application of agent-based modelling computer simulation that differs from types familiar to political scientists – even to those who may have some knowledge of ABM. A simulation, whether computerised or not, is a model, and like any model, it is a rendering of a situation, phenomenon or process that leaves out much more of the 'bloomin' buzzin' confusion of the world' than it incorporates. However, by doing so models render the world they reconstruct open to disciplined analysis.

Computer models are exactly like any other model, but they are formal models, not verbal models. For complex verbal models to be run on a computer, they must first be rendered as computer code translations. If the translations are not precise, unambiguous and fully consistent, the model will not run. Unlike closed-form formal models (algebraically solvable micro-economic, rational choice or game theoretic models), which are precise, unambiguous, fully consistent, but fundamentally simple, computer models can be quite complicated, involving large numbers of 'players' (agents) as well as large numbers of dimensions.

To be sure, most political scientists using ABM techniques have followed Robert Axelrod's advice to Keep It Simple Stupid (KISS). Because of the large number of interacting agents, these models are still enormously more complex than algebraically tractable formal models, but nonetheless they are simple. The ethos Axelrod encouraged was to generate elegant outcomes from extremely simple algorithmic processes as a way to exploit the promise of computer-assisted ABM for solving highly abstract, generic problems that could not be well formulated in a simple two-person game or that could not be solved with reference to average behaviour across large populations.

In their analysis of types of ABM works, Lustick and Miodownik (2009) refer to this most simple variant under the heading 'Abstractions'. Given a large array of standardised and similar interacting units, an abstract modelling research design seeks to understand how varying the algorithms that produce behaviour of individual units, or how varying the number and distribution of types of these units, affects the behaviour of the entire array. Such work is not designed as a stylisation of a specific problem or even of a real kind of problem as it appears in the world, which is populated by highly organised, differentiated and distinctive structures – people, organisations, resources and expectations. Instead, abstractly designed models have as their target the behaviours of the elements from which portions of the real world of interest are constructed. The kind of general insights that can be produced by such abstract ABM work, although incapable of explaining outcomes in any particular case or even in any particular kind of case, can inform, illuminate, enrich and simplify our understanding of the real world by showing that very simple relationships at the micro-level can account for complex patterns at the macro-level. Because of their generality, the insights derived from this kind of ABM design, much like those generated from abstract game theory, are blunt but readily transportable as patterns of recognised relationships across the borders of many scientific domains and problem sets.

A good example of an abstraction study is Axelrod's (1997) paper on 'The Dissemination of Culture: A Model with Local Convergence and Global Polarization',[5] where he explores the determinants of the spread of cultural traits and how different degrees of homogeneity can appear in a society despite a great deal of latent heterogeneity. Using a basic social influence model of tendencies to conform among interacting individuals, Axelrod finds that the number of distinctive traits activated by agents in stable homogeneous regions decreases with increases in the size of the repertoire of traits each agent has, but increases with the overall number of traits present in the society. Here, the abstract relationship being investigated is the effect of interacting size of agent repertoire and the size of the overall pool of cultural traits available in the society for the variability of traits actually exhibited. An added finding in the study is linked to Axelrod's manipulation of the size of the population, regardless of the traits that define its complexion. He finds that increasing the size of the population (not the pool of traits within it) beyond a certain point *decreases* the heterogeneity of that population as measured by the number of different distinctive and stable regions of cultural activation. By studying the relationship between adjustments at the micro-level and emergent properties at the macro-level across a large population of agents, Axelrod leverages an abstract research design to suggest underlying reasons for the unpredictability of individual state boundaries and the conditions under which theories predicting cultural homogenisation as a result of increased communication are more likely to be correct.

Despite the prominence of the 'KISS principle' as advocated by Axelrod, and its reinforcement of agent-based modeller attention on the abstraction type of simulation, there is in ABM work, as in virtually all other domains, no 'one best way' to design research. Much depends on the problem addressed and the ambitions and purposes of the researcher. While some researchers may have a strong taste for the elegance and power achievable by studying the emergent effects of simple abstractions, others may be attracted to wider, more complex horizons. Clearly, experiments focused on relationships among simple variables nested within very general theories (as the concept of size of identity repertoire is nested within constructivist identity theory) can generate data to produce powerful insights with wide ranging applicability. For researchers with suitably abstract tastes, the elegance and scope of such findings can provide much that is satisfying. Yet, such studies could only rarely contribute directly to those debates within political science or other social sciences that are animated by specific theoretical claims and counterclaims about dynamics within social fields featuring both complex agents and distinctive institutions.

As understanding of the dynamics and implications of specific simple mechanisms in general theoretical contexts increases, it is natural for researchers to ask more complex and nuanced questions. In just this way, it is natural for agent-based modellers to think about how more specific but still generic problems can be modelled. By building theoretically stylised models of familiar types of empirical domains, research designs can be used to examine competing hypotheses that have been deployed to explain patterns in those domains. For instance: one can design models of culturally divided states to examine their susceptibility to secessionism; model polities that may (or may not) be liable to experience national mobilisation; simulate the conditions under which international systems may be more or less disposed to stability; explore authoritarian regimes' vulnerability to latent challengers and more. This type of research design we label 'ensemble'.

The intention of an ensemble model is to create a replicable simulation of a specific *kind* of problem. The building blocks of an ensemble are not simple algorithmic mechanisms, though they may be included in the model, but operationalisations of theoretical positions with enough empirical corroboration to attract the interest of the researcher. In this type of research, design rules governing agent interactions are explicitly guided by good (i.e. empirically corroborated), or at least interesting theories in ways that are consistent with their premises. This creates multiple analytic opportunities including identification of inconsistent suppositions, exploration of the logical implications of interactions among theories and location of the range of possible and likely variation in outcomes in the domain of interest.

The objective of ensemble modelling is not only to translate the independent variables of the theories into formal and dynamic simulations, but also to register outcomes in the virtual world that are parallel to the outcomes predicted in the real world as stylised by the assumptions of the competing theories. Cederman (1997) uses this strategy in his book *Emergent Actors in World Politics* utilising evolutionary and constructivist theory to explore constitutive processes of international relations and critique the radical separation prevailing in neorealist international relations theory between processes of nation-state emergence and nation-state behaviour. The book is organised as a series of ensemble-type agent-based modelling experiments, each of which is based on the disaggregation of states by regions, competing identity groups and formulae, leadership and economic resources, different kinds of military technologies and problematically established state boundaries.

One such study focuses on 'nationalist coordination', involving an ensemble model of the process by which a divided state can become more or less consolidated around a unifying national identity. This model featured interacting mechanisms expressing predominance by a centre over a divided periphery, the distribution of resources among peripheral regions, degrees of overlap in identity repertoires across regions within the space of the state and the contribution of entrepreneurial leadership or its absence. He found that the 'more powerful the center compared to the periphery, the more inclusive the emergent identity of the new political unit', but that this relationship is strongly affected by the resource balance and intensity of competition across rival regions (Cederman 1997, p. 185). Cederman's study of nationalist coordination is not a study of the trajectory of any one particular case of nationalist integration or disintegration, so real-world data are not used for his model's initialisation. Instead, as an ensemble model, it represents a composite of real cases stylised to include only the abstract conditions identified as important for modelling purposes. Nor are Cederman's model outputs imagined to correspond to the past or future of any particular country or nationalist movement.

In contrast to abstraction or ensemble models, virtualisation studies entail simulations infused with data from a specific place and time to establish initial conditions. Outputs of the model are deemed to correspond to the possible futures of the target of the model. Virtualisation models will still be incomplete, that is, they will still be models, not parts of the world itself. Just as the mental model of a problem in an expert's head – for example, a question pertaining to the effect of a tsunami on ethnic conflict in Sumatra – would be an incomplete 'operationalisation' of the real world and of relevant aspects of it, so an ABM virtualisation of the problem capable of posing that question and generating possible trajectories into the future will not be an exact replica of the world. But in contrast to the expert's mental model, a computerised virtualisation is exact in its incompleteness, and it is replicable. Moreover, to the extent that the expert's mental model can be specified, and to the extent that those specifications can be translated into algorithms implementable in a virtual world, the model thereby created can serve as the basis for controlled experiments about implications of variation in assumptions, theoretical relationships, parameter values or initial conditions for the probability of different kinds of futures. Undoubtedly, ABM virtualisation research design is demanding, but it is not more demanding than the actual (if usually unacknowledged) tasks required of an expert asked to make a counterfactually based judgement about the effect of a policy option, or a specific set of circumstances or a historical institutionalist judging the causal implications of a historical antecedent for an observed state of affairs.

Using the Political Science-Identity (PS-I) agent-based modelling platform, Lustick and his collaborators have produced ensemble models to study secessionism in ethnically and regionally divided states (Lustick, Miodownik, and Eidelson 2004), the effects of violence and war on regime policies and prospects in the Middle East (Lustick 2004) and patterns of change in semi-authoritarian regimes (Lustick and Cartrite 2005). Using these models and techniques developed to build and experiment with them, and taking advantage of research and development support flowing from Defense Advanced Research Projects Agency (DARPA) through the Advanced Technology Laboratories wing of Lockheed Martin in a project called Integrated Crisis Early Warning System (ICEWS), we have developed virtualisation models of an assortment of Southeast Asian countries.[6] A subset of these models, for Malaysia, Indonesia, Vietnam, Sri Lanka, Bangladesh, Thailand and the Philippines, has been employed to retrodict outcomes in the early years of this century based on data available from the 1990s and to predict outcomes based on data available in 2008 and 2009 for 2009 and 2010. The trajectories which are the basis for these forecasts, that is, the distribution of sample runs for each country model, are also employed, heuristically, to identify mechanisms giving rise to observed outcomes and to think about answers to 'what if' or 'what would likely have happened if' type questions.

We do not here report on details of the performance of these models. However, as a basis for considering them as worthy of serious consideration and improvement, we note the following: using independently established 'ground truth' data with respect to the occurrence of 'rebellion' in Sri Lanka, Thailand, Bangladesh and Vietnam, our models of these countries tuned only with data available prior to the first quarter of the forecast year yielded probabilities of occurrence that departed from ground truth by an average of 0.55 on a scale anchored at 0 for 100% correct and 4.0 for 100% incorrect.[7]

One aspect of the models that produced these results is that they required the modellers to specify, as part of each model's instantiation, which groups in the country would be eligible for participating in rebellions and in which regions. That meant using expert knowledge to identify where there was at least a possibility of a rebellion and by whom – a rebellion defined as 'organised opposition whose objective is to seek autonomy or independence'. An important challenge has been to endogenise the emergence of rebellious groups so that within any particular trajectory of the model, PS-I can determine which groups in which regions are rebelling without modellers stipulating anything about this *ex ante*. For this purpose, and to systematise different mobilisation opportunities for all groups in the society with respect to a variety of events of interest – not just rebellion – we developed a module incorporated into each of the virtual country models. We call this module the dynamic political hierarchy (DPH). To illuminate the nature of ABM virtualisation modelling, we use the balance of this article to explain how the DPH operates and to illustrate its results by reporting on 2010 forecasts made with Virtual Thailand (VirThai) in an exercise of *prediction* rather than *retrodiction*.

The DPH is a set of routines implemented in PS-I that endogenises the nature of political action by agents affiliated with different groups. A detailed account of the DPH is presented below. As a first approximation, and to orient the reader in relation to the theories we shall discuss as undergirding its design, we summarise the operation of the DPH as follows. First, the DPH must identify and monitor changing relationships among groups of agents based on the degree of overlap of affiliations in agent repertoires. Each agent has a set of attributes which may contain more than one 'identity' (politically relevant affiliation). Each pair of identities in the repertoire of an agent represents an 'edge' in the network of groups that helps link the two groups marked by those identities. The DPH module then uses that information to update relations among groups to reflect changing propensities of different groups to challenge the political *status quo* with different intensities and different commitments to legality. Accordingly, the DPH combines (1) an architecture for networks of groups (sets of agents with a particular identity in their repertoire) to form and re-form based on amounts and patterns of overlapping membership in specific groups; and (2) a protocol enabling groups within the network to establish themselves in a hierarchy of involvement in and loyalty to the legal order.

2. Three theories and the DPH

The design of the DPH is intended as an operationalisation of the theories of cross-cutting cleavages, nested institutions and dynamic loyalties, combining them into a mechanism that regulates the intensity of political conflict based on the presence or absence of overlapping affiliations.

2.1. *Cross-cutting cleavage theory*

A basic element of the DPH architecture is the concept of society as a network of groups whose conflicts with one another can be mapped as aligned with or oblique to other groups. At one extreme, at any point in time, an absolute cleavage exists between two identities if no agent repertoires contain both identities. At the other extreme, one identity is perfectly aligned with a second identity if 100% of the agents with the first identity in their repertoires also have the second identity in their repertoires. Of course, this relationship need not be symmetrical. If the second identity is present in the repertoire of more agents with repertoires that do not contain the first identity, the second identity would not be *as* aligned with the first, as the first is with the second.

The concept of society as a network of overlapping identity affiliations, or cross-cutting cleavages, and the effect these cleavages have on politics has a rich history in the social sciences (Simmel 1955; Dahrendorf 1959; Lipset 1959; Dahl 1963; Allardt and Littunen 1964; Taylor and Rae 1969). Cross-cutting cleavages can be thought of as 'the interdependence of antagonistic groups and the crisscrossing within such societies of conflicts, which serve to "sew the social system together" by canceling each other out, thus preventing disintegration along one primary line of cleavage' (Coser 1965). A recent study by Dunning and Harrison (2010) adds to a body of the literature that illuminates the ability of cross-cutting cleavages to alter individual political behaviour. Dunning and Harrison analyse how a system of 'cousinage', a non-ethnic system of family alliances with origins in the thirteenth century, can help to explain the 'the limited presence of ethnic voting in Mali'. Typically, voters view candidates with whom they do not share an ethnicity with disfavour, but if that candidate shares a cousinage tie with the voter, the lack of coethnicity between the voter and the candidate is largely ignored in voting decisions.

2.2. *Nested institutions*

The nested institutions component of the DPH treats the structure of political authority as a layered array of modes of competition for power based on the scope of demands by different groups and expectations of compliance with norms and legally promulgated rules. This line of thinking began with Aristotle's theories about regimes of different kinds arising within cities (communities or polities) of different kinds and featuring incumbent leaders of different qualities and pathologies. It is a mainstay of the structural functionalist tradition in sociology and in systems theory.

Among political scientists, this tradition found its most influential expression in the work of David Easton. Easton's (1957) vision of political institutionalisation divides society into three levels: those who accept the political authority of the community, but not the regime – the legal order established within it; those who accept the regime's authority, but not that of the government established according to the laws of the regime and those who accept the government's authority. A fourth level can be considered to represent those in society who do not recognise, even in principle, the authority of the community. Easton's government level is comprised of those members of society who are ready to support the actions of the government taken in order to settle conflicts and handle conflicting demands. The regime level represents those who may or may not readily support the actions of the government, but who support the 'rules of the game'. These regime-level supporters may disagree on specific actions taken in order to address conflict, but they agree on the 'arrangements that regulate the way in which demands [on the government] are put into effect' (Easton 1957, p. 392). Even more broadly, community-level supporters are not concerned with the exact demands being put on the government or the framework within which demands of the government are made, but instead depend on being 'sufficiently oriented toward each other to want to contribute their collective energies toward pacific settlement of their varying demands' (Easton 1957, p. 391).

2.3. *Dynamic loyalties*

Work by Finifter and Linz provides the basis for the dynamics of loyalty dimension of the DPH. Finifter (1970) defined 'political powerlessness' as 'an individual's feeling that he cannot affect the actions of the government, that the authoritative allocation of values for the society which is at the heart of the political process, is not subject to his influence'. In Finifter's treatment, feelings of alienation can change, mainly as a function of changing norms, which lead groups and individuals to imagine themselves as more or less isolated, more or less powerful. She states that 'to the extent that social norms are dynamic . . . a state of isolation at one point in time may entail a different attitude set than isolation at another time'. It follows that political alienation, in addition to political isolation, is a dynamic state of affairs. A group that is virtually powerless today is not necessarily permanently powerless, and indeed these groups are often the source of populist movements and revolutions that completely shift the balance of power.

Linz shifted attention from feelings of powerlessness or alienation, per se, to different strategies of contestation among groups with different and dynamic levels of alienation or loyalty. He suggested a model more complex than the powerful vs. the powerless, or the elite vs. the mass, but of a hierarchy of propensities among groups to comply with established political authority in their mode of opposition – from 'loyal', to 'semi-loyal', to 'disloyal'. This framework was used to analyse patterns of stability and change associated with challenges to democratic regimes in the 1930s in Europe and re-democratisation in Latin America. Loyal groups are attached enough to the legal order to compete for political power only within the rules laid down and enforced by the government. Disloyal groups explicitly and completely reject the regime and are committed to extra-legal challenges to the institutions of the regime. Semi-Loyal groups are prepared to engage in a double discourse: behaving tactically as if loyal to the regime, but also sponsoring, serially or simultaneously, extra-legal activities to undermine the regime and take power illegally (Linz 1978).

The theoretical traditions used to guide both design and implementation of PS-I country models are expressed in non-formal, verbal and hence inexact forms. To be implemented for a computer-assisted ABM, these theories need to be formalised, that is rendered in a form that can operate with perfect replicability in the absence of human discretion. Accordingly, the tenets and assumptions of the models have to find unambiguous and complete and consistent specifications. The process of translating theory into a computer program is not only challenging, but also revealing. By forcing users to be complete and explicit, the process tends to highlight previously unappreciated theoretical ambiguity. If formal and complete theories were available, their instantiation in a computer program could be said to be achievable via a technique of 'derivation' – routines would be derived from theoretical propositions. But given the substantive limits of all social science theories when applied to complex problems and given the ambiguities present in most, the most we can say is that the construction of a computer model from routines is 'guided' by the theories we seek to implement or test.[8]

3. The dynamics of the DPH

The DPH is a vehicle for leveraging the explanatory power of the theories used to construct it, thereby integrating linear and non-linear shifts in the political landscape of a country. The balance of this article is devoted to explicating the operation of the DPH as guided by those theories – cross-cutting cleavages, nested institutions and dynamic loyalties.

PS-I country models combine a 'state' with a 'society'. These two fundamental elements of a country – the organisation of power and the array of motivated and resourced players interacting with or affected by that organisation – are established with a standardised ensemble-type model – the generic political model (GPM). The GPM is a template we use for the production of country-level virtualisation models. It is an ensemble-type PS-I model whose routines implement an ontology and a synthesis of theories and folk theorems about how a political system works. By adding the DPH module to the GPM, we were able to solve a significant modelling problem. Previous versions of our country models required us to predetermine regions where certain kinds of events, such as rebellions, could occur or stipulate, *a priori*, the groups that could and could not cause these events. By endogenising protocols determining the loyalty, semi-loyalty and disloyalty of groups at different points in time, the DPH has allowed us to overcome this restriction. In principle, rebellions, for example, can occur in any region where the markers of that activity appear and can be produced by the activity of any group.

Other modules within the GPM are guided by theories of the causes of anger, different types of mobilisation, the effect of violence and several other mechanisms for political change. For example, the theory of the effects of violence in our model, rather than referring to 'damage', refers to a drastic increase of negatives values at stake, which allows for an operationalisation that does not require decisions by the user about the physical damage of an event; only its social implications (see Lustick 2006). DPH categories are used in combination with the routines operationalising these theories to regulate the manner in which discontented agents mobilise. These behaviours (legal or illegal, non-violent or violent, discriminant or indiscriminant) thereby arise naturally within the course of each individual simulation run, helping produce different patterns of contestation, different outcomes and distinctive trajectories through the space of possible futures for the country model.

In view of space limitations which constrain our explication of the exact routines involved, we proceed in the following section to a nonetheless detailed account of how the DPH in our virtual country models operates to combine the intuitions these theories offer about relationships between changing cleavage patterns, differently structured networks of groups and different patterns of political mobilisation.

At every time step in the model, the DPH module takes a census of each agent's affiliations. Where overlaps between different groups exist, the DPH aggregates this information and then calculates how closely aligned each group is with every other group. The group affiliation publicly espoused most prominently is established as the 'Dominant' group. 'Incumbent' groups are identified as a separate category, as those groups comprised predominantly of agents aligned with the Dominant group via overlapping subscribed identities (affiliations included in their repertoires). Groups comprised predominantly of agents aligned with Incumbent groups (but not the Dominant group) are registered as 'Regime' groups. Groups comprised predominantly of agents aligned with Regime groups (but not with the Dominant or Incumbent group) are registered as 'System' groups. The DPH deems agents not registered in any of these groups as 'Non-system'.

In fact, the rules are somewhat more specific than this. The Dominant group consists of the group manifesting, at any given point in time, the largest aggregated influence of activated agents. The aggregated influence of a set of agents is calculated based on the influence levels of each agent in the set. The influence level of an agent reflects the relative impact of that agent's behaviour on other agents monitoring that behaviour. This influence level may range from very low (1) to very high (5). With an influence level of (1), a key aspect of the agent's behaviour, namely its activated affiliation, has only a marginal influence on agents monitoring that behaviour. With an influence level of (5), the agent's publicly espoused affiliation often decisively affects the behaviour of monitoring agents.[9]

Incumbent groups: This category consists of the groups of which it is the case that 70% of the agents in the group have the Dominant group identity in their repertoires. In other words, considering the set of agents that have identity X in their repertoire, namely set A_x, there is a subset of these agents that have the identity of the Dominant group in their repertoire, namely set $A_{x\mathrm{D}}$. If the number of agents that have both identities is 70% of the total number of agents with this identity in their repertoire, i.e. if $|A_{x\mathrm{D}}| \geq (70\% \times |A_x|)$, then group X is in the Incumbent category.

Regime groups: Similarly, at least 70% of the agents in a group in this category must have within their repertoires the identity of an Incumbent group or of the Dominant group (though the number of agents containing the Dominant identity in their repertoires must be less than 70%).
System groups: At least 70% of the agents in a group in this category must have within their repertoires the identity of the Regime groups, Incumbent groups, or the Dominant group (though the number of agents containing Incumbent or Dominant identities in their repertoires must be less than 70%).
Non-system groups: This category consists of isolated groups. In an isolated group, no more than 70% of agents can share identities with the Regime, Incumbent, or Dominant groups.

Two additional rules are implemented in the DPH to recognise the role of '*Staatsvolks*' and the implications of extreme poverty or powerlessness on cleavage overlap. A *Staatsvolk* is traditionally thought of as a group with tight and constitutive ties to the political order. The DPH treats any group whose identity appears in the repertoires of at least 50% of activated Dominant group agents as a *Staatsvolk* and prevents it from being categorised below the Regime level. A second rule governs groups experiencing extremely disadvantageous conditions (as produced by exogenously generated perturbations of the entire model). While experiencing these conditions, such groups must maintain a 90 rather than 70% link with a DPH category to be included within the category below. The effect of this rule is to reduce the integrative force of shared identities when one of the groups involved is perceived as powerless, impoverished or otherwise severely disadvantaged.

The DPH not only categorises groups and agents, but also systematically changes each agent's potential behaviour. Guided by Linz's theory of political legitimacy and of different strategies of contestation among groups, our model governs an agent's action according to its level in the DPH. If an agent is part of (subscribed to) the Dominant or an Incumbent group, it is attached enough to the legal order to contest other groups only within the rules laid down and enforced by the government. This mobilisation is referred to as lobbying. An agent that is not part of an Incumbent or Dominant group but is part of a Regime group may behave as if loyal to the Regime, but could also engage in illegal or legally ambiguous activities, referred to as protests. Lastly, agents outside the Dominant, Incumbent, or Regime category express discontent through violence.

4. The DPH and Thailand forecasts for 2010–2011

In part because of threshold effects associated with DPH transition rules that can translate small fluctuations into the re-categorisation of entire groups, fundamentally different political structures can emerge from essentially identical initialisations. To illustrate how the endogeneity of political processes contributes to the production of possible futures, then to the identification of plausible futures, and then to forecast of probable futures, we will draw on work done with our VirThai model. Specifically, we will use the DPH to discuss distributions of futures generated with this model and the forecasts they enabled for the period from August 2010 through July 2011.

It should be emphasised that none of the possible futures generated with the VirThai model can be inconsistent with theories operationalised in the model or with data used to initialise it. However, just because the model's trajectory has established that thread into the future as possible, does not establish it as plausible, let alone probable. To facilitate analysis and forecast, we narrow our focus from a sampling of possible futures to a sampling of plausible futures. We term the process by which the distribution of futures under consideration is narrowed from the set of possible futures to the set of plausible futures 'whittling'. We start with a large-*N* sample of model futures, generally 1000. We use observed 'ground truth' data about what has happened recently in a given country (in this case Thailand) to cull the subset of futures that diverge widely from this ground truth. For this experiment, the model was set to begin in November 2009. Accordingly, ground truth from October 2009 was used to whittle away a subset of the *possible* futures that, by November, could no longer be considered *plausible*. It was undoubtedly the case that in October 2009 the Red Shirt group was not the anchor of the dominant coalition in Thailand (not the Dominant group), we therefore removed from the 1000 futures produced by the VirThai model those that featured Red Shirt Dominance in October 2009.

Our procedure for replenishing the population is as follows. Once a subset of the futures is removed from the original 1000, the survivors are used to replace them. These replacements are generated as 'clones'. Each initialisation clone is then set off along distinctive paths, diverging from the trajectories from which they were cloned, by exposing it, over time, to a randomly generated but distinctive stream of exogenous shocks. This process of reintroducing the possible (via new streams of randomised perturbations) then narrowing the sample again to the plausible is used on a monthly basis beginning in November 2009 and continued through July 2010. Thus, each of the 1000 futures in our August 2010 distribution has survived a series of eight comparisons with ground truth in Thailand before continuing its trajectory from August 2010 to July 2011.

Among these 1000 trajectories of VirThai, five 'families', or types of futures, emerged as particularly prominent in the entire distribution. They are listed here in order of prominence in model weeks registered by these simulations with VirThai.

(1) *Buddhist Dominance* – (47.6%). The first family of futures, characterised most distinctly by the fact that the Buddhists are Dominant, features a very large number of Incumbent identities. In fact, the only two identities that are commonly found in the Regime, System or Non-system levels are Malay and Muslim.
(2) *Thai Ethnic Dominance* – (24.1%). The second family of futures, when the ethnic Thais are Dominant, is also characterised by a high number of Incumbent identities. However, in addition to Malay and Muslim, smaller ethnic groups, most notably the Isan, often fall into the Regime, System or Non-system level.
(3) *Yellow Shirt Dominance* – (19.1%). The third family of futures, when the Yellow Shirts are Dominant, contains few, if any, Incumbent identities. The Buddhists and Ethnic Thais are *Staatsvolk* groups in the Regime level, and virtually all the other major groups in Thailand, most notably the Red Shirts, the Malays, the Muslims and the ethnic Isan are found in the System or Non-system levels.
(4) *Thai Civic Dominant* – (8.6%). The fourth family of futures, when the Thai Civic group is Dominant, features the Red Shirts and the Isan as Incumbent identities while Buddhists and Ethnic Thais are both *Staatsvolk* groups in the Regime level. The Yellow Shirts are also found in the Regime level. Malays and Muslims are the only two groups that are commonly below the Regime level, found in the System level.
(5) *Red Shirt Dominance* – (0.5%). The fifth family of futures, when the Red Shirts are Dominant, also contain few, if any, Incumbent identities. The Buddhists and Ethnic Thais are again *Staatsvolk* groups in the Regime level and the Yellow Shirts join nearly all the other major groups in the System or Non-system levels.

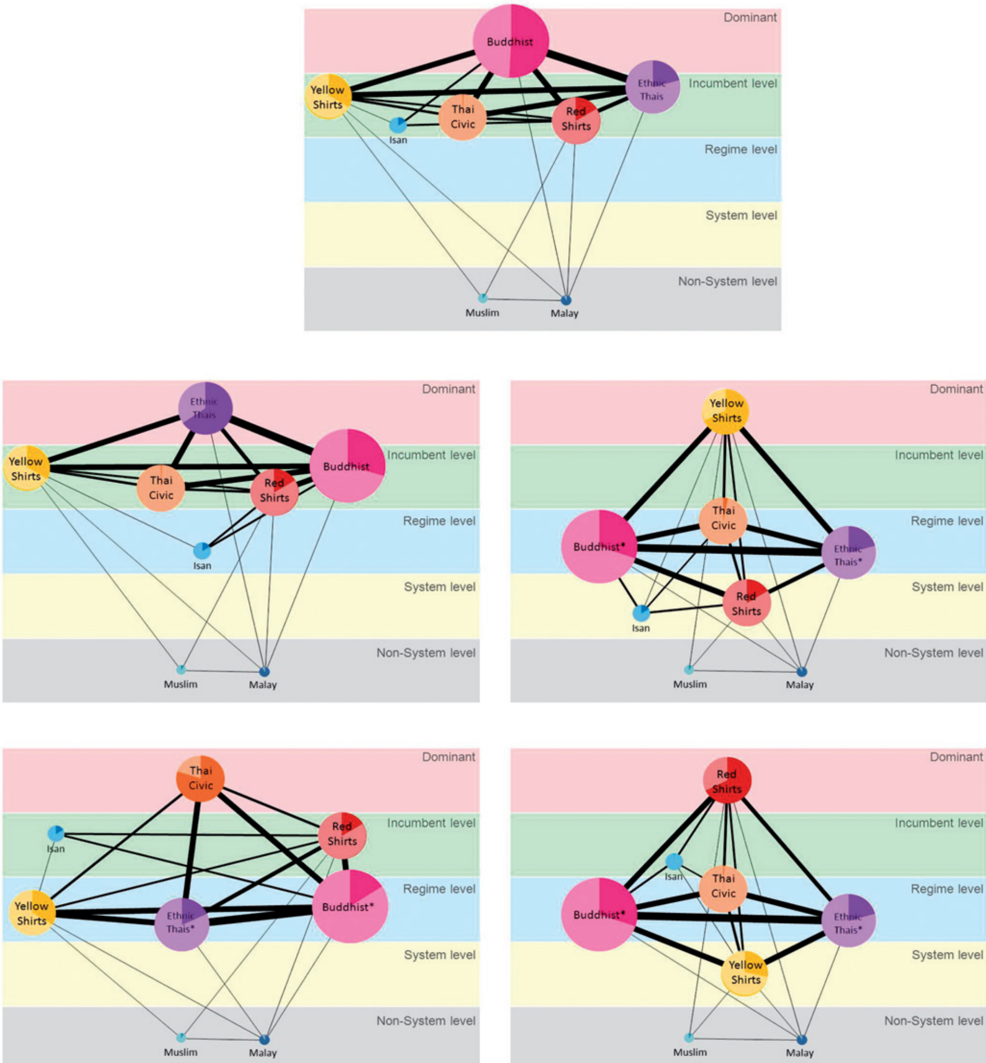

Figure 1. Representative schematics for five DPH families.

Figure 1 includes schematic representations of actual model outcomes for each of the five DPH families outlined above. In these schematic representations, each circle represents a group in Thailand and is labelled accordingly. The vertical location of the centre of each circle denotes where in the DPH that group is located (horizontal placement within a category is not instructive). The size of the circle represents the size of the group (how many agents are subscribed to the identity). The darker shaded portion of each circle represents the portion of the group currently mobilised (how many agents are activated on the group identity). The thickness of the line between two circles represents the strength of the connection between the groups (how many agents are subscribed to both identities). An asterisk (*) designates a group's categorisation as a function of its status as *Staatsvolk*.

Two of these DPH families, featuring Buddhist and Thai Ethnic groups as Dominant, include more than two thirds of the predicted time periods in our 1000 futures. In the depictions of these DPH configurations included in Figure 1, we can see that because so many inhabitants of Thailand are either Buddhist or ethnically Thai, when either of those groups is Dominant, most other significant groups in the society are allied with the government and therefore occupy the Incumbent category. By contrast, the Yellow Shirt Dominant, Thai Civic and Red Shirt Dominant families all feature governments anchored in smaller groups, leading much more often to less unifying patterns of cross-cutting cleavages with significant groups operating at the Regime level, including Buddhist and Thai Ethnic. Interestingly, the instability-prone configuration of political power in actual Thailand for most of the last 2 years has been Yellow Shirt Dominant, i.e. matching a salient pattern, but not one of the two most prominent patterns in our distribution. It is also worth noting that in our distribution of 2010–2011 futures, the Malay and Muslim groups, who have been embroiled for many years in chronic disputes over governance and resource distribution in the 'Deep South' of the country, remain alienated from and antagonistic to the Thai political system, regardless of how it is organised.

In line with the objectives of the ICEWS challenge, 'events of interest' we have sought to forecast with the VirThai model have included rebellion, insurgency, domestic political crisis and ethnic/religious violence. These are political outcomes powerfully shaped by the way power and loyalty is institutionalised in a society and by the behaviour of different groups, groups with differently aligned interests and different propensities to accept the authority of governments constituted in different ways. Agents in the model mobilise and petition for change in different ways depending on where groups with which they are affiliated are located within the DPH. Different configurations of the DPH can therefore be expected to be associated with different patterns of outcomes we observe for these events of interest.

As explained previously, Dominant and Incumbent level agents lobby, Regime level agents protest and System and Non-system level agents either attack or become violent. As a result, we find that the Buddhist Dominant and Thai Dominant families are relatively stable states of the Thai political landscape. With all of the largest groups occupying the Incumbent level, petitions for change and political challenges virtually all occur within the legal framework and take the form of what we refer to as 'lobbying' or peaceable and legal petition activity and the normal exertion of political pressure. To be sure, in the case of the Thai Dominant family, the Isan ethnic group is categorised at the semi-loyal 'Regime level'. It is therefore liable to pose some threat of 'protests' that can straddle the line between legal and illegal mobilisation, but the overall likelihood of politically destabilising unrest or a domestic political crisis arising from Isan protests alone is very low.

By contrast, the VirThai model exhibits substantial instability when the DPH configuration belongs to either the Yellow Shirt or Red Shirt Dominant families. In each case, the Dominant group lacks support by weighty Incumbent groups and must contend with the potential for large-scale protests (based on the size of the groups occupying the Regime category) but also with the illegal and possibly violent opposition of their arch rivals – Red Shirts if the Yellow Shirts are in power and Yellow Shirts when the Red Shirts are in power. The only slight asymmetry between these two families is that members of the Isan ethnic group in northeastern Thailand, who tend to be impoverished and naturally allied with the rural and poor base of the Red Shirt movement, commonly occupy the Incumbent category when the Red Shirts are Dominant. The same is true of splinter parties formed after the dissolution of the large populist Thai Rak Thai in 2007.

The Thai Civic Dominant family is particularly interesting. Here the model produces a significant number of futures where a large but not massive group of agents sharing an affiliation with Thailand as a state for its citizens attracts Incumbent support from both the disadvantaged Isan minority and the class-based Red Shirt movement. It appears that when the model exhibits Thai Civic Dominance, high rates of Thai Civic activation occur throughout the northern rural areas. The Buddhists and Thai Ethnics are significantly affected by the high rates of political mobilisation by Thai Ethnic Buddhists who are also members of the Thai Civic group. This renders the Buddhists and the Thai Ethnics as *Staatsvolk* groups who accordingly do not express discontent by challenging the regime directly. Meanwhile the Yellow Shirt movement, maintaining its strength in the Bangkok region, avoids challenging the regime because although its arch Red Shirt rival is supporting the government, it is not anchoring the governing coalition.

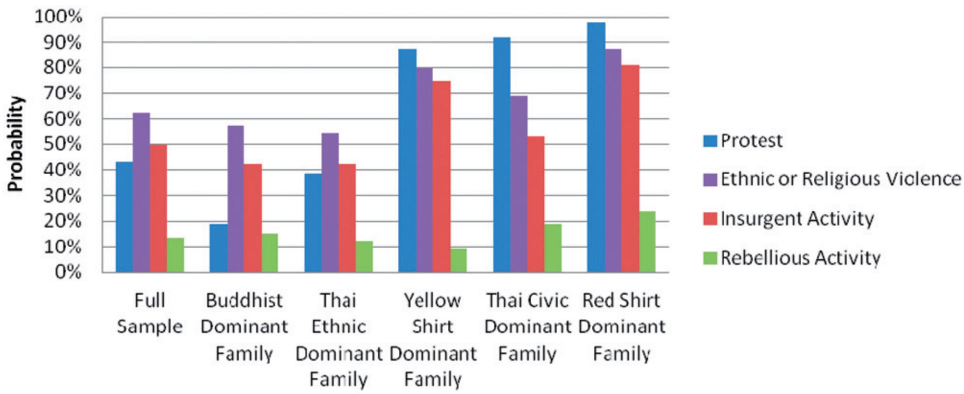

Figure 2. Chance of protest, ethnic or religious violence, insurgent activity and rebellious activity by DPH family.

The most straightforward forecasts made with the VirThai model for the year August 2010 through July 2011 yielded probabilities of events of interest (as defined within the overall ICEWS program) as occurring sometime during the year somewhere in Thailand. Rebellion, defined as organised opposition whose objective is to seek autonomy or independence, was forecast with a 73% probability. Insurgency, defined as organised opposition whose objective is to overthrow the central government, was forecast with a 100% probability. Domestic political crisis, defined as significant opposition to the government, but not to the level of rebellion or insurgency (e.g. power struggle between two political factions involving disruptive strikes or violent clashes between supporters), was forecast with a probability of 65%. Ethnic or religious violence, defined as violence between ethnic or religious groups that is not specifically directed against the government, was forecast with a probability of 91%.

The data reported in Figure 2 are different. It indicates the probability that the model will exhibit behaviour associated with the events of interest (though not necessarily the event itself) somewhere within Thailand at any one point in time in the 1000 model years of this experiment.

These data (Figure 2) are reported in a way designed to highlight the effects that the different shapes of the five DPH families have on the prevalence within our distribution of futures of protest, ethnic or religious violence, insurgent activity and rebellious activity. For example, the results reported in the 'Full Sample' column indicate that in 1000 model years, 'Insurgent Activity' appeared in just short of 50% of all time steps (with each time step representing approximately 1 week). On the other hand, during those periods when the Buddhist Dominant DPH configuration was present, the probability of any time step exhibiting Insurgent Activity was just over 40%.

We observe that each form of instability is most prominent within the (small) set of futures featuring Red Shirt Dominance. Although the likelihood of protest activity associated with the Thai Civic Dominance configuration is on par with that of Yellow Shirt and Red Shirt Dominance configurations, the prominence of violent efforts to overthrow the regime (insurgent activity) is considerably reduced, while rebellious activity (associated mainly with separatist actions by Malays in the Deep South) is higher than is associated with the Yellow Shirt Dominant configuration, but lower than is associated with Red Shirt Dominance. In the model, these patterns result from polarisation between Red and Yellow Shirt groups which pushes one or the other into an anti-regime posture if the other is Dominant, thereby setting the stage for angry agents in the out-of-power group to target the regime with violence. With respect to patterns of separatist rebellion or violent discontent in the Deep South, it appears that our model reproduces what is a rough consensus among experts on Thailand – that Monarchist/Yellow Shirt policies of patronage and cooptation towards discontented Malays tend to damp down levels of unrest in the Deep South (without eliminating them), in contrast to the more repression and exclusion oriented approach adopted by Red Shirt supported governments which tends to provoke escalations in the violence that chronically afflicts the region.

Our third and final probe into the VirThai model and its forecasts for 2010–2011 is to trace the main changes over time in one of the 1000 future trajectories. Above, we illustrated the different configurations of the DPH prominent in the population of VirThai futures by selecting one time-step profile from five different futures. Here, we choose a single future featuring a change from one DPH family to another and examine it at three different points in time, thereby highlighting the impossibility of understanding a dynamic model by examining a snapshot of it at any one point in time.[10]

The first image in Figure 3 is the DPH chart for time step 50 of the future we have selected for illustration purposes. At this time, the DPH is part of the Buddhist Dominant family with a concentration of large groups in the Incumbent level creating a relatively flat hierarchy. Note that the Buddhists are highly mobilised, as indicated by the relatively large area of the shaded portion of the Buddhist circle. The size of the Buddhist group and its effective rate of mobilisation allows the group to maintain its position as the anchor of the governing coalition dominance despite a similarly high level of mobilisation by the Yellow Shirts.

In the next image, the model has advanced to time step 75 of the same future. Buddhist is still the Dominant identity and its broad-based socio-political coalition is still intact. But Buddhist mobilisation has dropped significantly while Yellow Shirt mobilisation has grown. By drilling deeper into the model, we could determine what caused these shifts, but for our purposes here we can surmise that perturbations in the exogenously determined stream of small exogenous shocks experienced by each VirThai model were such as to incentivise a significant number of Buddhist Yellow Shirts to change the nature of their activation (mobilisation status) from Buddhist to Yellow Shirt.

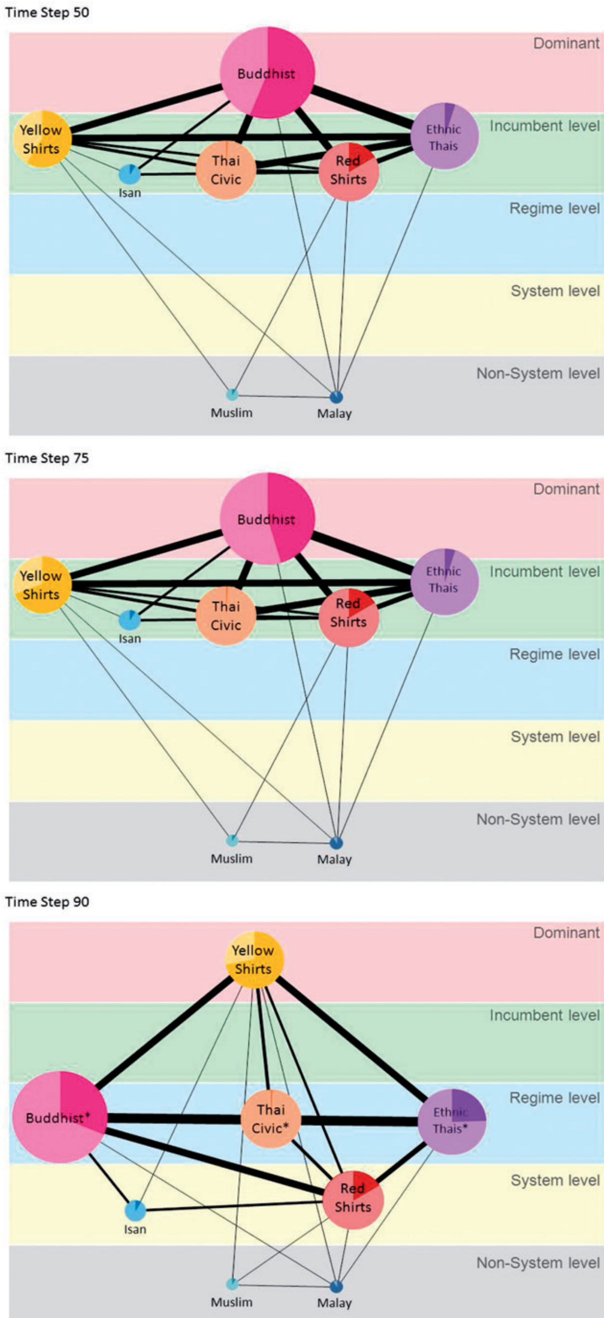

Figure 3. Yellow Shirt dominance emerging in a sample run.

By time step 90, depicted in the final image, Buddhist mobilisation has fallen even further allowing the Yellow Shirts to become the Dominant identity. The entire shape of the DPH changes accordingly, moving from a broad-based coalition to a politically isolated Yellow Shirt government facing the likelihood of violent unrest from severely alienated Red Shirts and Isan. Based on our analysis of the implications of different DPH configurations, we can expect at time step 90 that Thailand in this future will experience a significant uptick in unrest, including violent oppositionism. Instead of violent discontent arising only in the Malay/Muslim inhabited Deep South, this future may well include violence from elements of the Red Shirts, along with continued violent instability by the Malays and the Muslims. The Yellow Shirts will have to maintain a very high level of mobilisation to protect their Dominant position since large regime category groups will be poised to exploit the disequilibration that violent unrest creates to displace the Yellow Shirts from their control of the government.

5. Conclusion

We will have to see how this year turns out before ‘scoring’ our forecasts. However, how should we react if some of our forecasts turn out to be substantially incorrect? Would that, per se, disconfirm or invalidate the model or the technique? What, for example, should we learn if in the coming year there is no Insurgent or Rebellious activity in Thailand? What should we learn if what we observe what the model deems to be very unlikely, e.g. emergence of a Red Shirt Dominant political configuration stabilises the country? Unfortunately, perhaps, we cannot learn anything decisively from such outcomes, since the specific path followed by Thailand in all its concreteness will be a thread of actuality that is best understood as drawn from a very large distribution. Our models seek to describe the nature of that distribution, but not which future will, in the event, be drawn from it.

In other words, the very nature of these models is to anticipate that a large variety of futures are possible, just more or less implausible or probable. This is not to say that what happens in the real world cannot or should not affect the model. For example, if Thailand experiences a future this coming year which in key respects cannot be found in our sample, or if events that do occur and that we have forecast cannot be interpreted in relation to the mechanisms present in the model that produced them, then specific grounds would be available for reconsidering aspects of the modelling approach, including the theories animating the model and the manner in which those theories are operationalised.

Beyond the challenge of external validation, which ultimately requires a large number of models making large numbers of predictions or retrodictions about different countries and different years, realisation of the potential for computer-assisted agent-based modelling means overcoming substantial internal validity challenges. Meeting this challenge will entail taking advantage of what is in principle the full transparency of virtualisation models. Although it is necessarily true that the code of a model run successfully on a computer is translatable unambiguously into human language, the complexity of the algorithms used to translate theory into routines for agent interaction and updating as well as for data ingestion and model behaviour interpretation, can make tracing the workings of the model extremely difficult. As noted, Axelrod’s KISS approach to ABM was offered with precisely this problem in mind. To protect transparency and confidence in internal model validity, Axelrod, and the tradition of abstract modelling he encouraged, sought to keep model complexity to a minimum.

Indeed, there is no question that ensemble and virtualisation models do pose serious internal validity challenges. But in light of the unavoidable complexity of the important problems we seek to address, researchers should not allow that challenge to block their way. Furthermore, as domain-specific substantive theories continue to improve, abstract models for those domains, produced without the benefit of categories and relationships we confidently understand to be involved, will be rendered unpersuasive and even irrelevant. The acceleration of available computing power and the spread of agent-based modelling expertise are two other factors that make it more or less inevitable that the complexity of ABMs will increase, making it vital to develop best practices for enabling evaluation of internal model validity despite complexity.

The study reported in this article is part of a larger effort to meet this challenge. However daunting the task, researchers involved in the effort to surface and validate the complex implications of multiple theories integrated for the solution of ensemble or virtualisation-level problems can take heart from the fact that it is not a new challenge. It is simply a traditional challenge now acknowledged and addressed rather than ignored. For it is manifestly the case that establishing the 'internal validity' of verbal models or the implicit congeries of models used by experts to think about complicated problems, is a significantly more difficult task than seeking to do so for a model whose elements are known to be consistent, identical from one application of the model to the next, and at least in principle fully legible. Undoubtedly, ABM virtualisation research instruments, such as VirThai, are difficult to design, validate and interpret. But that these tasks are experienced as difficult is mainly attached to the fact that they enable such stringent demands to be made upon them. While such tasks are just as warranted in connection with the individual application of expert knowledge, it is difficult to even imagine how to subject inferences, forecasts and explanations arising from that source to the same rigorous examination.

It is typical of advances in any scientific domain that progress in the building of theory and in substantive understanding is often enabled by new investigatory tools, and that subsequent advances will also depend on further development of ancillary measurement and observational techniques. In our work, and in the ICEWS project as a whole, we have become convinced that techniques for visualising large amounts of data, especially multi-dimensional time-series data, have not yet matched the demand for these techniques that computer simulation modelling requires. In this article, we have made substantial efforts to render our models transparent, to help the reader 'see' what is going on in them and to track key aspects of the behaviour they exhibit and the data they produce. Yet, we do not consider these efforts to have been fully successful. We have been constrained, of course, by space limitations, but there are also difficult technological challenges for visualisation experts that ABM models pose and that deserve resources and urgent attention. We are confident that as those challenges are met, both analysis of data and replication of results will benefit in ways that realise the enormous potential of theory-guided, computer-assisted ABM simulations for investigating the social world.

Acknowledgements

This article was prepared for presentation at the Annual Meeting of the American Political Science Association, Washington, DC, September 2010. The research reported in this article was supported by the Defense Advanced Research Projects Agency via Lockheed Martin, ATL. Researchers interested in replicating findings should contact the authors for access to the necessary platform, templates and protocols.

Notes

1. For a discussion of the Shell Oil case and for a good summary of the technique of scenario writing and its contemporary antecedents, see Chermack and Walton (2006).
2. For excellent data and analysis documenting the distortive impact of scenario thinking on political forecasting efforts by those otherwise most likely to succeed at the task, see Tetlock (2005). See also below.
3. For a related, but distinct, argument emphasising the difficulties of prediction based on the predominance in human affairs of 'fat-tailed distributions', and the chaotic effects they entail, compared to normal distributions, see Mandelbrot and Hudson (2006).
4. BOGSAT: Bunch of Guys Sitting Around a Table.
5. Originally published by Axelrod (1997).
6. For a general account of this effort, known as the 'ICEWS', see O'Brien (2010).
7. For each country-year, each model was required to register a forecast probability of at least 67% to be credited with a correct forecast of either an occurrence or a non-occurrence.
8. The same is true of researchers seeking to test theories using aggregate data. Their problem of theoretical translation into testable models faces the double difficulty of underspecified theories and data that are only available in particular forms. To cope with these issues, researchers must respecify theoretical propositions into precise categories and then treat available data as if they were collected in exactly the categories specified by the theory. So, for example, theories purporting to explain the cessation of civil wars are typically tested by a definition of civil war as entailing 1000 battle deaths per year, or some such figure. Of course, the verbal theory never specified war in this way, but the apparatus of statistics requires the theory being applied to be 'guided' towards a specification that itself cannot be fully justified in terms of the theory.
9. Technically, the *aggregate influence* of a set of agents is calculated as follows:
   A_i = the set of agents activated on identity i;
   C_{ix} = the set of agents activated on identity i that have an influence level of x and
   $|C_{ix}|$ = the size of the set C_{ix}; in other words, the number of agents activated on identity i that have an influence level of x.
   Since there are five influence levels, then there are five sets $C_{ix} = C_{i1}$, C_{i2}, C_{i3}, C_{i4} and C_{i5}.
   In this way, the set A_i may be partitioned into influence levels such that
   $A_i = C_{i1} \cup C_{i2} \cup C_{i3} \cup C_{i4} \cup C_{i5}$
   The aggregate influence is calculated by multiplying the size of each of these sets by their associated influence level, and then adding these resulting values together.

$$\text{Aggregated influence of } A_i = (|C_{i1}| \times 1) + (|C_{i2}| \times 2) + (|C_{i3}| \times 3) + (|C_{i4}| \times 4) + (|C_{i5}| \times 5) = \sum_{x=1}^{5} (|C_{ix}| \times x).$$

10. One purpose of this exercise is to make it clear that the data reported on DPH configurations are quantified as model time steps featuring a particular configuration, with there being 60,000 model time steps in a population of 1000 year-long trajectories.

References

Allardt, E., and Littunen, Y. (1964), *Cleavages, Ideologies, and Party Systems; Contributions to Comparative Political Sociology*, Helsinki: Academic Bookstore.

Axelrod, R. (1997), 'The Dissemination of Culture: A Model with Local Convergence and Global Polarization', *Journal of Conflict Resolution*, 41, 203–226.

Cederman, L.-E. (1997), *Emergent Actors in World Politics: How States and Nations Develop and Dissolve*, Princeton, NJ: Princeton University Press.

Chermack, T.J., and Walton, J.S. (2006), 'Scenario Planning as a Development and Change Intervention', *International Journal of Agile Systems and Management*, 1, 46–59.

Coser, L.A. (1965), *The Functions of Social Conflict*, London: Routledge & Kegan Paul.
Dahl, R.A. (1963), *A Preface to Democratic Theory*, Chicago, IL: Chicago University Press.
Dahrendorf, R. (1959), *Class and Class Conflict in Industrial Society*, London: Routledge.
Dunning, T., and Harrison, L. (2010), 'Cross-cutting Cleavages and Ethnic Voting: An Experimental Study of Cousinage in Mali', *The American Political Science Review*, 104, 21–39.
Easton, D. (1957), 'An Approach to the Analysis of Political Systems', *World Politics*, 9, 383–400.
Finifter, A. (1970), 'Dimensions of Political Alienation', *The American Political Science Review*, 64, 389–410.
Linz, J.J. (1978), *The Breakdown of Democratic Regimes: Crisis, Breakdown, and Reequilibration*, Baltimore: The Johns Hopkins University Press, pp. 27–38.
Lipset, S.M. (1959), 'Some Social Requisites of Democracy: Economic Development and Political Legitimacy', *The American Political Science Review*, 53, 69–105.
Lustick, I.S. (2004), 'Simulating the Effects of Israeli-Palestinian Violence, Fundamentalist Mobilization, and Regional Disruption on Regime Stability and USA-Friendly Outcomes in Middle East Polity', in *Paper Prepared for Presentation at the Santa Fe Institute Workshop on Origins and Patterns of Political Violence 1: Violence in Civil Wars*, 16–18 January, Santa Fe, NM.
Lustick, I.S. (2006), 'Defining Violence: A Plausibility Probe Using Agent-based Modeling', in *Paper Presented at PRIO Workshop, The Role of First Actors in Civil Wars*, 17–18 August, Oslo, Norway.
Lustick, I.S., and Cartrite, B. (2005), 'Virtualstan: An Agent-based Modeling Strategy for the Comparative Dynamics of Authoritarian Regimes: Bureaucratic Authoritarianism, Bureaucratic Feudalism, and Neopatrimonialism', in *Paper Presented at the Annual Meeting of the American Political Science Association*, 1–4 September, Marriott Wardman Park, Omni Shoreham, Washington Hilton, Washington, DC.
Lustick, I.S., and Miodownik, D. (2009), 'Abstractions, Ensembles, and Virtualizations: Simplicity and Complexity in Agent-based Modeling', *Comparative Politics*, 41, 223–244.
Lustick, I.S., Miodownik, D., and Eidelson, R.J. (2004), 'Secessionism in Multicultural States: Does Sharing Power Prevent or Encourage It?', *American Political Science Review*, 98, 209–230.
Mandelbrot, B.B., and Hudson, R.L. (2006), *The (Mis)behavior of Markets: A Fractal View of Risk, Ruin, and Reward*, New York: Basic Books.
O'Brien, S.P. (2010), 'Crisis Early Warning and Decision Support: Contemporary Approaches and Thoughts on Future Research', *International Studies Review*, 12, 87–104.
Parker, G., and Tetlock, P.E. (2006), 'Counterfactual History: Its Advocates, Its Critics and Its Uses', in *Unmaking the West: 'What-If?' Scenarios that Rewrite World History*, eds. R.N. Lebow and G. Parker, Ann Arbor, MI: University of Michigan Press, pp. 363–392.
Simmel, G. (1955), *Conflict and the Web of Group Affiliations* (Transl. by K.H. Wolf and R. Bendix), Glencoe, IL: Free Press.
Taleb, N. (2007), *The Black Swan: The Impact of the Highly Improbable*, New York: Random House.
Taylor, M., and Rae, D. (1969), 'An Analysis of Crosscutting Between Political Cleavages', *Comparative Politics*, 1, 534–547.
Tetlock, P.E. (1999), 'Theory-driven Reasoning about Possible Pasts and Probable Futures: Are We Prisoners of Our Preconceptions?', *American Journal of Political Science*, 43, 335–366.
Tetlock, P.E. (2005), *Expert Political Judgment: How Good Is It? How Can We Know?*, Princeton, NJ: Princeton University Press, pp. 189–215.
Tetlock, P.E., and Lebow, R.N. (2001), 'Poking Counterfactual Holes in Covering Laws: Cognitive Styles and Historical Reasoning', *American Political Science Review*, 95, 829–843.
Tetlock, P.E., and Parker, G. (2006), 'Counterfactual Thought Experiments: Why We Can't Live with Them and How We Must Learn to Live With Them', in *Unmaking The West: 'What If?' Scenarios that Rewrite World History*, eds. P.E. Tetlock, R.N. Lebow and G. Parker, Ann Arbor, MI: The University of Michigan Press, pp. 14–44.