Choice, Decision, and Measurement: Essays in Honor of R. Duncan Luce

Edited by

A. A. J. Marley

LEA
LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
1997 Mahwah, New Jersey London

Choice, Decision, and Measurement
A. A. J. Marley (Ed.),

A STOCHASTIC MODEL FOR THE EVOLUTION OF PREFERENCES

J.-CL. FALMAGNE, M. REGENWETTER, AND B. GROFMAN

University of California, Irvine

ABSTRACT. This chapter presents a model describing the evolution of preferences as a stochastic process. These preferences are represented by weak orders, i.e. rankings with possible ties, on a set of alternatives, and can be modified under the influence of 'tokens' of information delivered by the environment according to a stochastic mechanism. The parameters of this mechanism can be estimated from the data and are descriptive of the environment. The potential effect of a token is to move an alternative up or down in an agent's ranking. Attitude change is modeled by the stepwise transitions between the weak orders, which takes the form of a Markov process. The model permits exact predictions (up to a small number of parameters) of panel data in which the judges have been required to repeatedly evaluate the alternatives at times $t_1, \ldots, t_n$. An illustrative application of this model is described in a companion paper (Regenwetter, Falmagne, & Grofman, 1995). That illustration uses NES Thermometer (Rating) data on the 1992 presidential candidates.

1. INTRODUCTION

A few quantitative models are available, in the political science and economics literature, that deal with attitude change (e.g., Anderson, 1971; Converse, 1964, 1975; Zaller, 1992).

Some of these models are intended to apply to panel data (Converse & Markus, 1979; Markus, 1982; McPhee, Andersen, & Milholland, 1962) where the same individuals in a large sample have been questioned repeatedly, at times $t_1, t_2, \ldots, t_n$. On each occasion, the panel members have been asked to express their preference concerning a fixed set of alternatives. While some of the published models have a probabilistic component (modeling for example the measurement error), they typically do not cast such data as a manifestation of a stochastic process in the specific sense of this term in the theory of stochastic processes (e.g., Parzen, 1962; Norman, 1972). As a result, the predictions, and the ensuing analysis of the data, are not as complete and revealing as they could be. In particular, we cannot compute, in terms of the parameters of these models, the joint probability of observing the preference relations $P_1, P_2, \ldots, P_n$ at times $t_1, t_2, \ldots, t_n$ (for any choice of n-tuples of preference relations and times of observation).

Key words and phrases. Attitude change, public opinion, preferences, strict weak order, Markov process.

Acknowledgments. This work was supported by NSF grant No SBR 9307423 to Jean-Claude Falmagne at the University of California, Irvine. We thank two referees for helpful comments.

Address for correspondence. J.-Cl. Falmagne, School of Social Sciences, University of California, Irvine CA 92717. Email: jcf@uci.edu

This paper offers a stochastic model of attitude change that is closely related to recent work in mathematical behavioral science (Falmagne, 1996; Falmagne & Doignon, in press). The key idea is consistent with social choice theory in that, at any time $t > 0$, an individual's attitude is represented by a preference relation on the set of alternatives. These preference relations are formalized by 'weak orders,' that is, rankings with possible ties. A current ranking may be altered by 'tokens' of information delivered by the environment in a probabilistic fashion. The potential effect of a token is to move an alternative up or down in the current ranking. The axioms specifying the model ensure that detailed predictions can be obtained. In particular, we can derive an exact expression for the asymptotic probability of any ranking $\succ$ (see Theorem 3). We can also compute the joint probability of observing rankings $\succ$ and $\succ'$ at time t and $t+\delta$, respectively, for any pair of rankings $\succ$ and $\succ'$ and for large t (see Theorem 4). Moreover, the model lets us explicitly solve for the degree of relative negativity/positivity in the information token environment (see Regenwetter et al., 1995), e.g., to evaluate the extent of negative campaigning by reasoning backwards from observed changes in voter preferences. Negative campaigning has been studied by, for instance, Garramone (1985), Skaperdas and Grofman (1995), Ansalobehere, Iyengar, Simon, and Valentino (1994). For concreteness, our attention will be focused on the special case of three alternatives. A followup paper (Regenwetter et al., 1995) provides an illustrative test of the model with National Election Study (NES) panel data on voter evaluations of political figures, with a focus on the Bush-Clinton-Perot 1992 contest. That paper also contains a detailed discussion of the statistical issues associated with the practical application of the model.

Sections 1 and 2 respectively contain an informal outline of the model, and a point by point comparison of its assumptions with other concepts and approaches. A formal statement of the assumptions is given in Section 3 in the guise of three Axioms. The predictions are listed and examined in Section 4. Section 5 discusses some limitations of the present model and possible elaborations. Section 6 contains a summary of the work, and some suggestions for further research.

2. Overview of the Model

We write $\mathcal{A}$ for the set of alternatives. By a *preference relation* on the set $\mathcal{A}$, we mean here a binary relation in the usual sense of set theory, i.e., a set of ordered pairs of elements of $\mathcal{A}$. We restrict consideration to those preference relations $\succ$ which satisfy the following condition: for all i, j and k in $\mathcal{A}$,

$$\text{[SW]} \qquad \text{if } i \succ j, \text{then} \quad \begin{cases} \text{not } j \succ i \text{ and} \\ \text{either } i \succ k \text{ or } k \succ j \text{ (or both).} \end{cases}$$

A preference relation satisfying this condition is sometimes called a *strict weak order* (Roberts, 1979), and we conform to this terminology. Condition [SW] induces a strict weak order on the set $\mathcal{A}$. In fact, it is well known that if the set $\mathcal{A}$ is finite or countable and the relation $\succ$ satisfies [SW], then we can assign a number $u(i)$ to each alternative i in such a way that

$$i \succ j \quad \text{if and only if} \quad u(i) > u(j) \tag{1}$$

(see e.g. Krantz et al., 1971). Note that the scale u is only defined up to an arbitrary strictly increasing transformation. The empty relation $\emptyset$ – i.e., the relation containing no ordered pairs – is a special case of a preference relation. Indeed, [SW] 'vacuously' holds in that case because the case $i \succ j$ never arises. In the case of a set $\mathcal{A} = \{1, 2, 3\}$, there are exactly 13 different strict weak orders on that set. They are represented by their graphs (Hasse diagrams) in the 13 rectangles of Figure 3. (Ignore the other features of that Figure for the moment. Note that the empty strict weak order is represented by the empty rectangle in the middle of the figure.) We only consider the case of three alternatives for the rest of this paper.

The model is organized around four basic ideas.

1. The latent strict weak orders. We suppose that the responses of an individual to some questions of a survey are governed by a latent personal strict weak order which we call the 'state' of that individual. In the case of three alternatives, there are thus 13 possible states. The empty strict weak order is referred to as the *neutral* state. The set of all states will be denoted by $\mathcal{S}$.

2. The naive state. We also assume that an individual is initially naive, in the sense that, when first confronted with the set $\mathcal{A}$, he or she does not prefer any alternative to any other one. In other words, the state of a naive individual is the neutral state.

3. The probabilistic environment. Starting from this initial state, successive transformations may take place over time. Specifically, we assume that the individual is immersed in a probabilistic environment delivering at certain random times $t_1, t_2, \ldots, t_n, \ldots$ 'tokens' of information regarding the alternatives. These tokens represent events occuring in the environment and having a positive or negative connotation regarding particular alternatives.

4. The tokens. Examples of token generators are T.V. programs, newspaper articles, or conversations with acquaintances convincingly extolling or criticizing some alternative (Iyengar & Kinder, 1987; Iyengar, Peters, & Kinder, 1982). In our model, however, the occurrence of the tokens is not regarded as part of the recorded data. They are not necessarily observable, or controllable by the social scientist. Nevertheless, these tokens play a crucial role in the mechanism responsible for the evolution of the preferences and deserve close attention. In fact, we shall see that the analysis of the data will reveal statistical aspects of the occurrence of the tokens, thereby providing useful indications regarding the information flow.

As indicated above, a token can be *positive* or *negative.* Moreover, to each of these two types of tokens corresponds its *opposite* token. We shall indulge in some idealization and gather all the positive tokens pertaining to alternative i into one class, which will be denoted by $[i]$. For convenience, we shall refer to that class as 'token $[i]$.' A similar idealization will apply to all token types. Table 1 summarizes this notation.

The potential effect of a positive token $[i]$ is to modify the current state so as to move alternative i to the top position of the strict weak order, that is, the position in which i dominates the two other alternatives. There are three such

Types of Tokens	*Representing Symbol*
alternative i is good	$[i]$
i is not (necessarily) good	$\widetilde{[i]}$
i is bad	$[-i]$
i is not (necessarily) bad	$\widetilde{[-i]}$

TABLE 1. The four types of tokens and their notation.

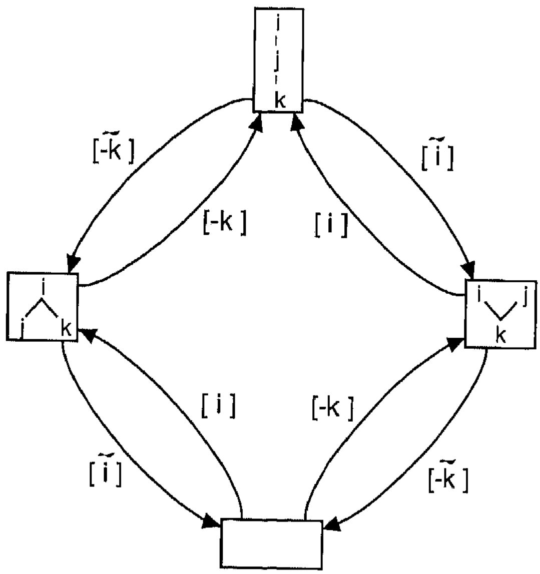

FIGURE 1. Effects of the tokens on the various states.

states, namely[1]:

$$[i \succ j, i \succ k], \quad [i \succ k \succ j], \text{ and } [i \succ j \succ k]. \tag{2}$$

(Note the convenient abuse of notation committed in Equation 2: since each of the three formulas specifies a different strict weak order, different symbols – e.g., $\succ$, $\succ'$ and $\succ''$ – should have been used. Our convention simplifies the writing and will be used whenever the context makes clear what is intended.) For example, an individual in state $[i \succ k, j \succ k]$ (that is, preferring both i and j to k) and perceiving token $[i]$ would end up in state $[i \succ j \succ k]$. This transition is represented by the upward arrow in the upper right corner of Figure 1. (The generic graph of Figure 1 illustrates all the possible transitions and will be useful for the rest of this Section.)

The opposite of token $[i]$ is denoted by $\widetilde{[i]}$. As indicated, the occurrence of such a token will modify the current state by removing alternative i from its top position

[1]The notation should be self explanatory. For example, $[i \succ j, i \succ k]$ means that i is prefered to both j and k, which are indifferent to each other, and $[i \succ k \succ j]$ has the same meaning, with moreover k prefered to j.

(if it occupies it). Two instances of such a transition are represented in Figure 1. The downward arrow in the upper right corner represents the transition which is the opposite of the one we just described. The downward arrow in the lower left corner represents the transition from the state $[i \succ j, i \succ k]$ to the neutral state. Note that the opposite tokens, when they are effective, always transform a state into one nearer the neutral state. Note also that a token is not always effective. The states are endowed with some rigidity, in the sense that transformations only take place between adjacent states. For example, the occurrence of token $[i]$ has no effect on state $[j \succ i, k \succ i]$ because each of the states having i in its top position (the three states in Equation 2) is far removed from $[j \succ i, k \succ i]$. The intuition underlying the assumption of partial rigidity of the states is that extensive changes in preferences cannot be triggered by a single token of information. (For example, a piece of negative information concerning an alternative at the top of someone's ranking is likely to be discarded.) This relation of 'adjacency' of the states could be formalized, but we shall refrain from doing so here. In any event, its meaning is made clear by the graph of Figure 1 in the case of three alternatives.

The effect of a token ζ on a state $\succ$ will be captured by an operation $\circ$ which is defined by the graph of Figure 1. Thus, the operation $\circ$ maps the pair $(\succ, \zeta)$ to some strict weak order $\succ' = \succ \circ \zeta$. For example:

$$\begin{aligned} [2 \succ 3 \succ 1] \circ \widetilde{[-1]} &= [2 \succ 3, 2 \succ 1] \\ [2 \succ 3 \succ 1] \circ \widetilde{[-3]} &= [2 \succ 3 \succ 1] \\ \emptyset \circ [-3] &= [2 \succ 3, 1 \succ 3]. \end{aligned}$$

One rationale for the transformations illustrated by Figure 1 is to imagine that the alternatives are implicitly evaluated by the respondents on a 3-point scale having -1, 0 and +1 as possible values, with 0 serving as a reference point. Equivalent alternatives are always rated 0. (This means that at most one alternative can be rated 1, or -1.) The value 1 corresponds to the top position of the state, when only one alternative occupies that position (cf. alternative i in the last two cases of Equation 2). A value -1 corresponds to the bottom position, with the same proviso. The value 0 given to some alternative j corresponds to the remaining three pairs of cases, namely:

$$\begin{array}{cc} [i \succ j \succ k], & [k \succ j \succ i], \\ [j \succ k, i \succ k], & [k \succ i, j \succ i], \\ [i \succ j, i \succ k], & [k \succ j, k \succ i]. \end{array}$$

The 0 position on the 3-point scale will sometimes be referred to as the *middle* position (even though j may be a maximal or a minimal element in the state, as in the last four examples above). It is easy to check that exactly 13 strict weak orders can be generated by these rules, via Equation 1. The effect of a positive token $[i]$ or a negative token $[-i]$ is to add or substract 1 to (from) the current value of alternative i if that value is currently 0. The corresponding opposite tokens $\widetilde{[-i]}$ and $\widetilde{[i]}$ have the reverse effect and set the current value of an alternative to 0 when it has been -1 or 1, respectively.

We shall exercise these concepts with an illustrative hypothetical sequence of tokens gradually transforming the states, starting from the neutral state. These transformations are represented in Figure 2. The time axis is in the first column, the horizontal bars marking the occurrence of the tokens, with time flowing from the top to the bottom of the figure. The tokens themselves are indicated in the second column. The current state is pictured in the third column by its Hasse diagram. At time 0, the agent is in the neutral state, and remains in that state until the occurrence of token [1] at time t_1. This results in transforming the neutral state $\emptyset$ into state

$$[1 \succ 2, 1 \succ 3] = \emptyset \circ [1].$$

Next comes the negative token $[-2]$ at time t_2. Its effect is to move alternative 2 to the bottom of the state, resulting in the state

$$[1 \succ 3 \succ 2] = [1 \succ 2, 1 \succ 3] \circ [-2].$$

Token $\widetilde{[1]}$ occurs at time t_3, displacing alternative [1] from its top position, and moving the state closer to the neutral state. Note that token [2], occuring at time t_4, has no effect on the current state $[1 \succ 2, 3 \succ 2]$. The reason is that changing that state into a state having alternative 2 in the top position would be a major transformation, which cannot be realized by a single token. We leave to the reader to ponder the effect of token $\widetilde{[-2]}$ occuring at time t_5.

At first blush, the opposite tokens may appear superfluous. Intuitively, it may perhaps seem that their role could be reassigned – via some appropriate modification of the rules of the model – to the positive and negative tokens. They were introduced to give the model some flexibility in capturing important effects. For example, a prevalence of opposite tokens in the environment would induce a high probability of the neutral state, in other words, a high proportion of uncommitted individuals. In particular, in a political context, a high proportion of opposite tokens of the types $\widetilde{[i]}$ or $\widetilde{[-i]}$ could yield a large number of uncommitted voters.

Finally, we suppose that the delivery of the tokens by the medium during the period of reference is, at least to a first approximation, a stable process in the sense that the probabilities of occurrence of the various tokens in any interval of time $[t, t+\delta]$, $t > 0$, $\delta > 0$, do not vary with t. (This idea will be specified in Section 3.) An important consequence of the Axioms is that the temporal succession of the states is a homogeneous random walk on the family of all states, with the probabilities of the transition between the states being governed by the probabilities of the tokens. A graph of this random walk is displayed in Figure 3.

To simplify the graph, only the centrifugal transitions – i.e., away from the neutral state – are indicated. The centripetal transitions can be obtained by reversing the arrows and capping each of the symbols representing tokens by a '$\sim$' sign. (The symbols β_i and τ_i refer to the transition probabilities between the states and should be ignored for the moment. They will be explained in Section 3.) One useful feature of this random walk is that asymptotic results can be obtained (see Section 3). In particular, it can be shown that the asymptotic probabilities of the

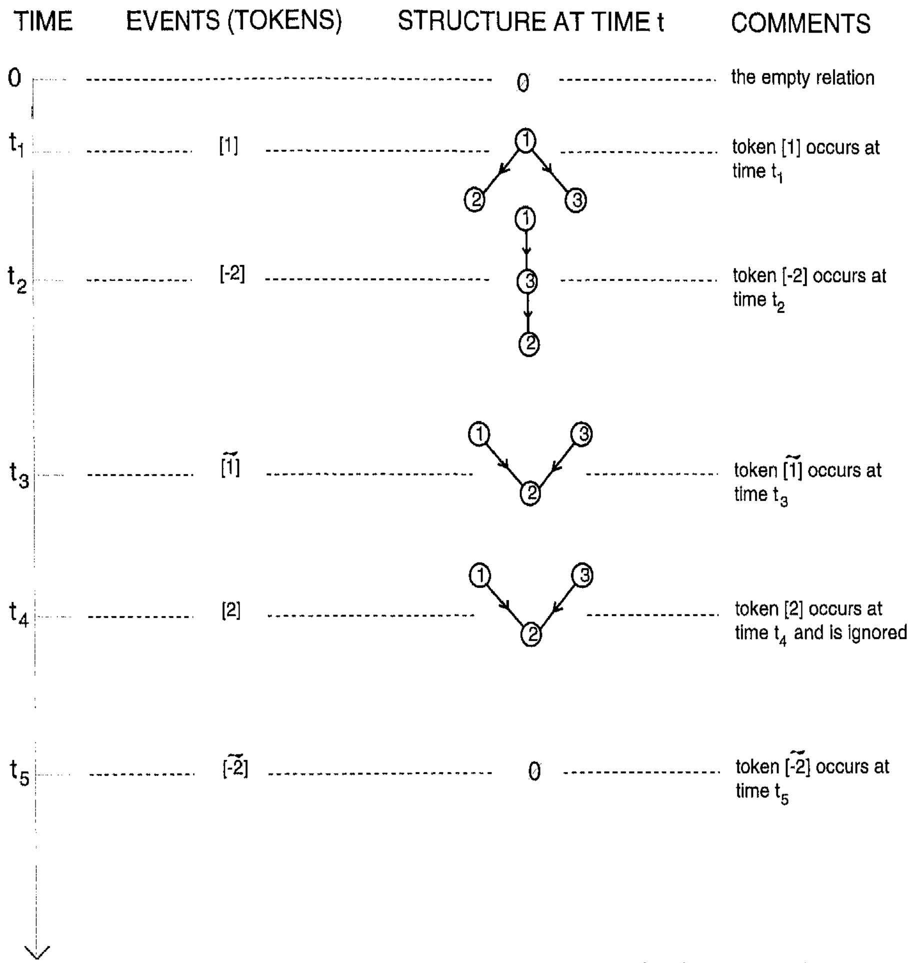

FIGURE 2. Illustrative hypothetical sequence of tokens occuring at times $t_1, \ldots, t_5, \ldots$ and the resulting preference relations.

states satisfy the following regularity condition. In words:

The asymptotic probabilities of two adjacent states $\succ$ and $\succ \circ \zeta$ differ by a factor which depends on the token ζ, but not on the state $\succ$. (3)

To state this condition more precisely, some notation is required. For any state $\succ$, let $p_\succ$ be the asymptotic probability of that state[2]. Take any two distinct states $\succ$ and $\succ'$ and let ζ be some token. Suppose that the four states

$$\succ, \quad \succ \circ \zeta, \quad \succ', \quad \succ' \circ \zeta$$

[2]In other words, suppose that the population of reference has been exposed to the tokens for a long time. Then, $p_\succ$ is the probability that, if an agent is sampled from the population, then this agent will be in state $\succ$.

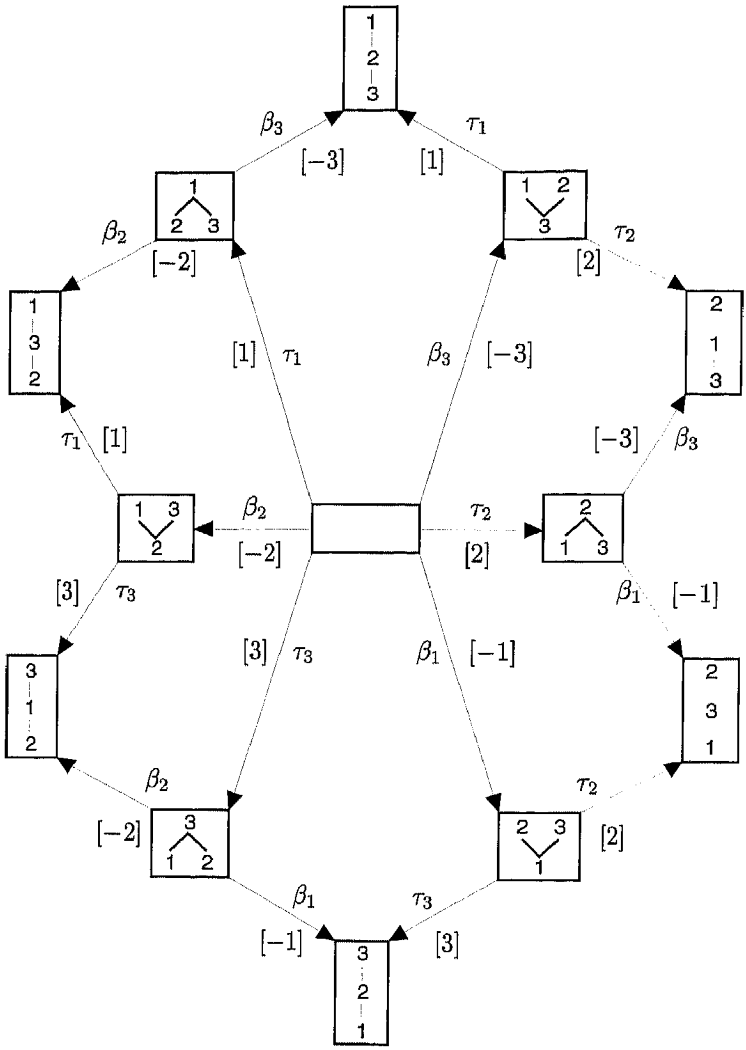

FIGURE 3. Transition diagram of the random walk on $\mathcal{S}$. The positive or negative token producing a transition is marked next to the corresponding edge, with its probability. To simplify the graph, only the centrifugal transitions (that is, away from the neutral state) are indicated. The centripetal transitions are obtained by reversing the arrows and by capping the symbols by '$\sim$' (tildes). This graph contains 6 different instances of the generic graph of Figure 1 which are obtained by setting i, j and k equal to 1, 2 and 3 in the 6 possible ways.

are all distinct. As a straightforward consequence of Theorem 3 (see Section 8 at the end of this chapter), we must have

$$\frac{p_{\succ}}{p_{\succ\circ\zeta}} = \frac{p_{\succ'}}{p_{\succ'\circ\zeta}}. \tag{4}$$

For example:

$$\frac{p_{\emptyset}}{p_{\emptyset\circ[-3]}} = \frac{p_{[1\succ 2,1\succ 3]}}{p_{[1\succ 2,1\succ 3]\circ[-3]}} = \frac{p_{[2\succ 1,2\succ 3]}}{p_{[2\succ 1,2\succ 3]\circ[-3]}},$$

that is,

$$\frac{p_{\emptyset}}{p_{[1\succ 3,2\succ 3]}} = \frac{p_{[1\succ 2,1\succ 3]}}{p_{[1\succ 2\succ 3]}} = \frac{p_{[2\succ 1,2\succ 3]}}{p_{[2\succ 1\succ 3]}}.$$

(cf. the three edges marked $[-3]$ at the upper right of Figure 3). Equation 4 means that the ratio $p_{\succ}/p_{\succ\circ\zeta}$ does not depend on the state $\succ$. Writing $G(\zeta)$ for the right member of Equation 4, and multiplying both members by $p_{\succ\circ\zeta}$, we get

$$p_{\succ} = p_{\succ\circ\zeta} \times G(\zeta)$$

which is a formalization of Equation 3. The invariance of ratios condition specified by Equation 4 is of interest because it has received some empirical support (see Regenwetter et al., 1995).

Another result in the same vein deals with two successive polls separated by a time interval δ. This result, which is too technical to state here (see Equation 14) concerns the joint probability $p_{\succ,\succ'}(\delta)$ of observing strict weak order $\succ$ at time t and strict weak order $\succ'$ at time $t+\delta$, for large t and for any pair of strict weak orders $\succ, \succ'$. This type of result provides a quantitative prediction of the effect of the passage of time on the correlation between successive judgments given by the same individuals.

An axiomatic presentation of this model is given in Section 3. It is worth pointing out that this model is only one of a fairly large class based on the same ideas. In these models, the successive preference relations of an individual are also seen as a realization of a stochastic process, and similar mechanisms manufacturing the changes are postulated. The models differ by the type of tokens considered, by the particular type of preference relations adopted for the states and by the transformations of the states induced by the tokens (i.e., the operation$\circ$ of this paper). We chose to present the strict weak order case in view of the good fit to the data obtained by Regenwetter et al. (1995). Other models, centered on different relations – linear orders in Falmagne (1996), and semiorders, partial orders and interval orders in Doignon and Falmagne (in press) (see also Doignon & Falmagne, in press) were less successful. Note that, even though all these models are conceptually related, the technical differences between them are not trivial.

3. Comparison with other models of attitude change

There are two prominent differences between the model that we propose and other models of attitude change, such as those encountered in the literature on survey research dealing with panel data (Converse, 1964), or in the social psychology and political science literatures on persuasion and attitude change[3] (Fishbein & Ajzen, 1981; Iyengar & Kinder, 1987; Zaller, 1992).

One concerns the type of preference relation considered. Measures of attitudes are often regarded as providing cardinal (i.e., numerical) information. For instance, thermometer data or Likert scales are treated as instances of interval scale measurement. A much more conservative viewpoint is taken here, in which only the order information contained in the numerical responses is retained: we suppose that an agent's attitude is represented by some strict weak order $\succ$, which may change over time. As indicated by Equation 1 such a relation can be given a numerical interpretation, but the numerical scale contains no cardinal information. (For example, the difference $u(i) - u(j)$ between two scale values has no empirical meaning.) However, the potential loss caused by discarding cardinal information is largely compensated by the comprehensive analysis of the data made possible by the model, which relates to the second of the two differences. In contrast to other models, attitude change in our model is explicitly cast as a (real time) stochastic process, and as such takes into account the full history of the responses provided by each agent. In the case of a sample of agents tested at times $t_1, t_2, \ldots, t_n$, these data correspond to the observed frequencies of all the n-tuples $(\succ_1, \succ_2, \ldots, \succ_n)$ of strict weak orders. For three alternatives, and with agents having provided their preferences on three occasions, the data comprises $13^3 - 1 = 2196$ independent frequencies, to be explained with only a dozen or so parameters.

We must also mention in passing the vast literature in psychology dealing with mathematical learning theory (Bush & Mosteller, 1955; Atkinson, Bower, & Crothers, 1965; Norman, 1972). There is a parentage between the learning models developed by the mathematical psychologists and that of this paper. In fact, there was an early attempt to apply mathematical learning theory to social choice theory (Suppes, 1961). The empirical context is quite different, however. The mathematical learning models apply primarily to highly controlled situations in experimental psychology. A representative example is a task in which a respondent is asked to predict which of two lights will come on. The respondent indicates his or her prediction, at the beginning of each trial, by pressing one of two keys. The trial ends up with a probabilistic event providing feedback. Three aspects of this approach are relevant here. For one, the occurrence of the stimuli (e.g., lights) is controlled by the experimenter (while the tokens are hypothetical constructs, having probabilities that have to be estimated from the data.) For another, with very few exceptions, the learning models are 'discrete parameter stochastic processes', that is, they are formalized in terms of a sequence of discrete trials (while the model of this paper formalizes a process taking place in real time, that is, a 'continuous parameter stochastic process'). Finally, the class of possible 'states of the subject' in the learning models does not resemble that postulated here, namely, a family of preference relations. These empirical and theoretical differences prevent a wholesale importation of the mathematical learning models.

[3]There is also a vast literature on propaganda and persuasion (Petty & Cacioppo, 1981; Hastie, 1986). With some adaption, the sequential features at the focus of that work could be incorporated in our approach.

4. Formal Statement of the Model

We consider three basic sets. We recall that $\mathcal{A} = \{1,2,3\}$ is the set of alternatives. Each of the 13 strict weak orders on $\mathcal{A}$ is a possible *state* of an individual, and we write $\mathcal{S}$ for the set all such states. Note that $\mathcal{S}$ contains the empty relation $\emptyset$. The set of all tokens is denoted by $\mathcal{T}$. Thus, $\mathcal{T}$ contains the positive and negative tokens and their respective opposites. We have

$$\mathcal{T} = \{\zeta | \zeta = [i] \text{ or } \zeta = \widetilde{[i]} \text{ or } \zeta = [-i] \text{ or } \zeta = \widetilde{[-i]}, \text{ for } i = 1,2,3\}.$$

The effect of a token ζ on a state $\succ$ will be captured by the operation $(\succ, \zeta) \mapsto \succ \circ \zeta$ already encountered and defined by the graph of Figure 1

We suppose that there exists a probability distribution $\theta : \zeta \mapsto \theta_\zeta$, with $\theta > 0$, on the set $\mathcal{T}$ of all tokens. This probability constrains the delivery of the tokens via a Poisson process governing their times of occurrence (see Axiom [T]). The model is cast in terms of three collections of random variables. We write:

$\boldsymbol{S}_t$ to signify the state of the individual at time $t > 0$,

$\boldsymbol{N}_{t,t+\delta}$ to specify the number of tokens arising in the half open interval of time $]t, t+\delta]$, with $t > 0$. In the sequel, we write $\boldsymbol{N}_t = \boldsymbol{N}_{0,t}$,

$\boldsymbol{T}_t$ to mean the last token presented before or at time t. We set $\boldsymbol{T}_t = 0$ if no tokens were presented, that is, if $\boldsymbol{N}_t = 0$.

Thus, $\boldsymbol{S}_t$ takes its values in the set $\mathcal{S}$ of states, $\boldsymbol{N}_{t,t+\delta}$ is a nonnegative integer, and $\boldsymbol{T}_t \in \mathcal{T} \cup \{0\}$. Notice that $\boldsymbol{N}_t$ will turn out to be the 'counting random variable' of a Poisson process, specifying the number of Poisson events occuring in the interval $]0, t]$. The three axioms below recursively define the stochastic process $(\boldsymbol{S}_t, \boldsymbol{N}_t, \boldsymbol{T}_t)$ in terms of the parameters θ_ζ and one parameter λ specifying the Poisson process governing the times of occurrence of the tokens.

Axioms

[I] (*Initial state.*) Initially, the state of the agent is the neutral state $\emptyset$. The agent remains in state $\emptyset$ until the realization of the first token. That is,

$$\mathbb{P}(\boldsymbol{S}_t = \emptyset | \boldsymbol{N}_t = 0) = 1\,.$$

The notation $\mathcal{E}_t$ stands for any arbitrarily chosen history of the process before time $t \geq 0$; $\mathcal{E}_0$ denotes the empty history.

[T] (*Occurrence of the tokens.*) The occurrence of the tokens is governed by a homogeneous Poisson process of intensity λ. When a Poisson event is realized, the token ζ occurs with probability θ_ζ, regardless of past events. Thus, for any nonnegative integer k, any real numbers $t \geq 0$ and $\delta > 0$, and any history $\mathcal{E}_t$,

$$\mathbb{P}(\boldsymbol{N}_{t,t+\delta} = k) = \frac{(\lambda\delta)^k e^{-\lambda\delta}}{k!} \tag{5}$$

$$\mathbb{P}(\boldsymbol{T}_{t+\delta} = \zeta | \boldsymbol{N}_{t,t+\delta} = 1, \mathcal{E}_t) = \mathbb{P}(\boldsymbol{T}_{t+\delta} = \zeta | \boldsymbol{N}_{t,t+\delta} = 1) = \theta_\zeta\,. \tag{6}$$

[L] (*Change of state.*) If an individual is in the state $\succ$ at time t, and a single token ζ arises between times t and $t+\delta$, then the individual will be in state $\succ \circ \zeta$ at time $t+\delta$, regardless of past events before time t. Formally:

$$\begin{aligned} & I\!P(\boldsymbol{S}_{t+\delta} = \succ' \,|\boldsymbol{T}_{t+\delta} = \zeta, \boldsymbol{N}_{t,t+\delta} = 1, \boldsymbol{S}_t = \succ, \boldsymbol{N}_t = k, \mathcal{E}_t) \\ &\qquad = I\!P(\boldsymbol{S}_{t+\delta} = \succ' \,|\boldsymbol{T}_{t+\delta} = \zeta, \boldsymbol{N}_{t,t+\delta} = 1, \boldsymbol{S}_t = \succ) \\ &\qquad = \begin{cases} 1 & \text{if } \succ' = \succ \circ \zeta, \\ 0 & \text{if } \succ' \neq \succ \circ \zeta. \end{cases} \end{aligned}$$

For the rest of this paper, we assume that these three Axioms hold.

Remarks. Three types of objections can be raised against these axioms. One concerns the stability inherent in the homogeneous Poisson process postulated in Axiom [T]. In the case of a political election, it may appear quite unrealistic to suppose that the intensity of the campaign does not significantly vary over time. Actually, it is well known that the opposite is true. The second objection is that, in the framework of the model, the only difference between the agents polled lies in the chance occurrence of the tokens, the distribution of which is the same for all agents. This is implausible, since the agents read different newspapers, watch different TV programs, live in different neighborhoods, all of which should reasonably affect the individual distributions of the tokens[4]. A third potential criticism is that the model requires that all the tokens are 'processed' by the agent, without allowing for possible selective perception/selective attention effects (Iyengar, 1990; Zaller, 1992).

All of these objections are well taken, but less critical than they may appear, because they only bear on superficial aspects of the theory. Relatively minor changes in the assumptions are easy to conceive, which would go a long way toward eliminating these shortcomings, leaving the basic machinery of the model intact. We return to some of these issues in Section 5.

5. Results

We first consider a particular realization of the Poisson process at times t_1, t_2, $\ldots$, t_n, $\ldots$. The ensuing sequence of states is a discrete stochastic process, which turns out to be a random walk. Let us state this formally.

We partition the time axis into the segments

$$]0, t_1[, \; [t_1, t_2[, \; \ldots, \; [t_n, t_{n+1}[, \; \ldots \tag{7}$$

such that $\boldsymbol{N}_t = 0$ for $t < t_1$, $\boldsymbol{N}_t = 1$ for $t_1 \leq t < t_2$, and in general $\boldsymbol{N}_t = n$ for $t_n \leq t < t_{n+1}$. Fixing the sequence (t_n), we define the discrete parameter process $\boldsymbol{S}^*_n = \boldsymbol{S}_{t_n}$ with state space $\mathcal{S}$. The process $(\boldsymbol{S}^*_n)$ is called the *discrete companion* of $\boldsymbol{S}_t$. Even though this discrete parameter process is implicitly indexed by the particular sequence of times of occurrence of Poisson events, in some important sense it does not depend on it. In fact, we have the following situation:

[4]Still more elaborate differences between agents have been considered (and modeled, for the case of two agents) by McKelvey and Ordeshook (1986).

Theorem 1. *With t_1, t_2, ... , t_{n+1} as in (7), the discrete companion $(\boldsymbol{S}_n^*)$ of $(\boldsymbol{S}_t)$ is a homogeneous random walk on the set $\mathcal{S}$ of all states, with transition probabilities defined, for any $\succ, \succ'$ in $\mathcal{S}$, by*

$$p^*_{\succ,\succ\circ\zeta} = I\!P(\boldsymbol{S}^*_{n+1} = \succ \circ\, \zeta | \boldsymbol{S}^*_n = \succ) = I\!P(\boldsymbol{S}_{t_{n+1}} = \succ \circ\, \zeta | \boldsymbol{S}_{t_n} = \succ) = \theta_\zeta.$$

This result follows immediately from the definitions. (All proofs are contained in Section 8. Proofs.) We call the Markov chain $(\boldsymbol{S}_n^*)$ a random walk because its transitions only take place between neighbor states, where $\succ' \neq \succ$ is a *neighbor* of $\succ$ if $\succ' = \succ \circ\, \zeta$ for some token ζ. The graph of this random walk is pictured in Figure 3. Note that only the centrifugal transitions and their probabilities are marked. A simplified mnemonic notation is used for these probabilities. We write:

$$\theta_\zeta = \begin{cases} \tau_i & \text{if } \zeta = [i] \\ \tilde{\tau}_i & \text{if } \zeta = \widetilde{[i]} \\ \beta_i & \text{if } \zeta = [-i] \\ \tilde{\beta}_i & \text{if } \zeta = \widetilde{[-i]}. \end{cases} \tag{8}$$

Thus, the movements to and from the top of the state are induced by tokens with probabilities τ_i and $\tilde{\tau}_i$, respectively, and the movements to and from the bottom of the state are induced by tokens with probabilities β_i and $\tilde{\beta}_i$, respectively.

We define by

$$p^*_{\succ,\succ'}(k) = I\!P(\boldsymbol{S}^*_{n+k} = \succ' | \boldsymbol{S}^*_n = \succ) \tag{9}$$

the k-step transition probability of the random walk $(\boldsymbol{S}_n^*)$. We also have:

Theorem 2. *The stochastic process $(\boldsymbol{S}_t)$ is a homogeneous Markov process, with transition probability function*

$$p_{\succ,\succ'}(\delta) = I\!P(\boldsymbol{S}_{t+\delta} = \succ' | \boldsymbol{S}_t = \succ) = \sum_{k=0}^{\infty} p^*_{\succ,\succ'}(k) \frac{(\lambda\delta)^k e^{-\lambda\delta}}{k!}. \tag{10}$$

The 'long term' probability of any state $\succ$ exists and can be computed. Specifically:

Theorem 3. *For any state $\succ$, we have the asymptotic probabilities*

$$p_\succ = \lim_{t\to\infty} I\!P(\boldsymbol{S}_t = \succ) = \lim_{n\to\infty} I\!P(\boldsymbol{S}^*_n = \succ) \tag{11}$$

$$= \frac{\xi_1(\succ)\xi_2(\succ)\xi_3(\succ)}{\sum_{\succ'\in\mathcal{S}} \xi_1(\succ')\xi_2(\succ')\xi_3(\succ')}, \tag{12}$$

with, for every $\succ \in \mathcal{S}$ and distinct $i, j, k \in \{1, 2, 3\}$,

$$\xi_i(\succ) = \begin{cases} \tau_i\tilde{\beta}_i & \textit{if } i \succ j \textit{ and } i \succ k, \\ \tilde{\tau}_i\beta_i & \textit{if } j \succ i \textit{ and } k \succ i, \\ \tilde{\tau}_i\tilde{\beta}_i & \textit{in the remaining cases.} \end{cases} \tag{13}$$

Writing K for the denominator of Equation 12, we have thus for the four generic cases:

$$p_{[i\succ j\succ k]} = \frac{1}{K}\tau_i\tilde{\beta}_i\tilde{\tau}_j\tilde{\beta}_j\tilde{\tau}_k\beta_k \tag{14}$$

$$p_{[i\succ j, i\succ k]} = \frac{1}{K}\tau_i\tilde{\beta}_i\tilde{\tau}_j\tilde{\beta}_j\tilde{\tau}_k\tilde{\beta}_k \tag{15}$$

$$p_{[j\succ i, k\succ i]} = \frac{1}{K}\tilde{\tau}_i\beta_i\tilde{\tau}_j\tilde{\beta}_j\tilde{\tau}_k\tilde{\beta}_k \tag{16}$$

$$p_{[\emptyset]} = \frac{1}{K}\tilde{\tau}_i\tilde{\beta}_i\tilde{\tau}_j\tilde{\beta}_j\tilde{\tau}_k\tilde{\beta}_k. \tag{17}$$

Remarks. 1) The conclusion of Theorem 3 in the guise of Equations 12-13 would hold under much more general assumptions on the process governing the delivery of the tokens. The homogeneity of the Poisson process plays no useful role in establishing the result. In fact, a general class of renewal counting process would do as well. On the other hand, this homogeneity is critical for the next theorem (for the second time interval, of duration δ).

2) The 13 asymptotic probabilities $p_\succ$ are expressed by Equation 12 in terms of 12 parameters: τ_i, $\tilde{\tau}_i$, β_i, $\tilde{\beta}_i$, $i = 1, 2, 3$, with a sum equal to 1. However, these probabilities only depend upon the probabilities of the tokens via the products $\xi_i(\succ)$, with $i \in \{1, 2, 3\}$ and $\succ \in \mathcal{S}$. There are nine such products. In fact, a still more economical reparametrization of the model is available, involving only 6 parameters (see Section 8. Proofs). Suppose that the frequencies of all 13 strict weak orders in a sample of respondents have been collected experimentally. This remark means that the $13 - 1$ independent frequencies of the strict weak orders can be predicted by 6 parameters. This kind of prediction, while not trivial if the number of respondents polled is large, is not breathtaking. The predictive power of the model is more impressive in the case of the sequential statistics to which we now turn.

Using Theorem 2 and Theorem 3, an explicit expression can be obtained for the joint probability of observing the states $\succ$ and $\succ'$ at time t and time $t + \delta$, respectively, for large t.

Theorem 4. *Let $p^*_{\succ,\succ'}(k)$ be as in Equation 9, i.e. the k-step transition probability between the states $\succ$ and $\succ'$ in the random walk $\mathbf{S}^*_n$ on $\mathcal{S}$. Then, successively*

$$\lim_{t\to\infty} \mathbb{P}(\mathbf{S}_t = \succ, \mathbf{S}_{t+\delta} = \succ') = \lim_{t\to\infty} [\mathbb{P}(\mathbf{S}_t = \succ) \cdot \mathbb{P}(\mathbf{S}_{t+\delta} = \succ' \mid \mathbf{S}_t = \succ)]$$

$$\begin{aligned} &= p_\succ \cdot p_{\succ,\succ'}(\delta) \\ &= \frac{\xi_1(\succ)\xi_2(\succ)\xi_3(\succ)}{\sum_{\succ''\in\mathcal{S}}\xi_1(\succ'')\xi_2(\succ'')\xi_3(\succ'')} \sum_{k=0}^{\infty} p^*_{\succ,\succ'}(k)\,\frac{(\lambda\delta)^k e^{-\lambda\delta}}{k!}, \end{aligned} \tag{18}$$

with $\xi_i(\succ)$ $(i \in \{1, 2, 3\}, \succ \in \mathcal{S})$ as in 13.

Remark. Suppose that, on two occasions separated by a time interval of δ days, the respondents in a sample have been asked to provide rankings of the three alternatives, allowing for possible ties. This means that $13^2 - 1 = 168$ independent frequencies of all the possible pairs of strict weak orders have been collected. Assuming that the first poll took place after the respondent had an extended exposure to the flow of tokens (that is, at a sufficiently large time t), Equation 18 provides a prediction of these 168 independent frequencies in terms of 12 parameters: 12-1 token probabilities, plus the parameter λ of the Poisson process[5]. In the case of three successive polls, we would have $13^3 - 1 = 2196$ independent frequencies predicted by the same number of parameters. Further discussion on statistical issues for models of this kind can be found in Doignon and Falmagne (in press) or Regenwetter et al. (1995).

6. Limitations of the Model and Possible Elaborations

Examining the assumptions of the model from a common sense standpoint leads to several possible criticisms. One concerns the homogeneity of the Poisson process governing the times of occurrence for the tokens. Consider the case of three political candidates running for an elective office. It is unrealistic to suppose that the intensity of the campaign – which is measured in the model by the density of the tokens – would not vary over time. A more sensible assumption is that campaigning starts slowly, then builds up to culminate just before the election. Remember however that some key predictions of the model do not depend upon this homogeneity assumption. In particular, the formula for the asymptotic probabilities $p_{\succ}$ of the states in Equation 12 would be valid for a general class of renewal counting processes, including non-homogeneous ones. On the other hand, Equation 18 in Theorem 4 critically depends upon the homogeneity of the Poisson process. What this means is that, from the viewpoint of these homogeneity considerations, the model as stated here would be appropriate for the following situation involving two polls. One poll taken just before a major event taking place long after the beginning of the campaign, such as a debate between the candidates, while the second poll is taken, a short time later, to assess the effect of the debate. Two different sets of token probabilities θ_ζ and θ'_ζ would be required. The first set of probabilities θ_ζ would be used for the computation of the asymptotic probabilities $p_{\succ}$ in the first factor of Equation 18, which does not depend upon homogeneity. The second set of probabilities θ'_ζ reflects the effect of the debate, and are used in the computation of the k-step transition probabilities $p^*_{\succ,\succ'}(k)$ in the summation part of Equation 18. The assumption of homogeneity is critical for the prediction of these transition probabilities, but not implausible since it only covers a short interval of time.

In any event, while the assumption of homogeneity greatly simplifies the computation of the predictions, it is by no means essential to the basic mechanisms driving the model. We could assume, for example, that the intensity parameter λ of the Poisson process is a function of time, the analytic form of which may depend upon a couple of parameters. In the case of the three political candidates mentioned above, the estimated value of these parameters would be revealing of the general course of the campaign.

[5]This supposes that the probabilities of the tokens do not vary for the period considered. If we assume that the probabilities of the tokens change, we can split the time period into two periods with different parameters. For example as a result of a political debate between the candidates that took place just after the first poll, a new process may start right after the first poll. Then the number of parameters increases to $6 + 11 + 1 = 18$.

Another possible criticism concerns the assumption of homogeneity of the population of individuals. Three objections must be distinguished here.

1. All the individuals are initially in the same (neutral) state.
2. They have the same chance of being exposed to the tokens.
3. All the individuals have the same reactions to the tokens.

It is easy to take care of Objection [1]. Rather than assuming that any individual is initially, with probability one, in the neutral state, we could postulate an a priori distribution on the set of states, which would be representative of the population of agents under consideration. This added touch of realism would change very little in the model, but would come at the cost of 12 (= 13 − 1) extra parameters. This would not be prohibitive, if the data consist in several polls. Objection [2] can be met, for example, by considering different subpopulations of individuals. The model can then be applied separately to the samples selected from each subpopulation. The estimated values of the probabilities θ_ζ of the tokens can be compared, and may reveal informative differences between the subpopulations. As for Objection [3], notice that the model is not exactly stating that the effect of a token is the same for all agents, since this effect depends on the current state. Nevertheless, it may be of interest to complicate the model by introducing an extra mechanism modulating the effect of a token that has occurred, thus modifying the transition probabilities between the states. This type of elaboration of the model would be in the spirit of the work of Sniderman, Glaser, and Griffin (1990), who allow for different categories of voters to use the same information differently, and the work of Iyengar (1990) that deals with selectivity biases in information monitoring.

There are various ways of developing our model in that direction. We shall outline two of them. For concreteness, suppose that the data allow one to sort the agents into two categories labeled R and D. In some situations, the effect of a token may depend on the category of the agent. To model such a dependency, we could index the operation $\circ$ by R or D. Thus, the effect of a token ζ on a individual in state $\succ$ and belonging to category D would be $\succ \circ_D \zeta$. In this version, the tokens are perceived the same way by all agents, but the definition of the operation $\circ$ depends upon the agent's category.

Another possibility involves the assumption of different screening mechanisms on the part of the agents. A token presented may be ignored with some probability depending on the category of the agent. From a formal viewpoint, this assumption is equivalent to supposing that the probability distribution on the tokens is not the same for the two categories of agents. Accordingly, only very minor changes of the model are required.

Further elaborations of the model can be gathered from the discussion in Falmagne and Doignon (in press). While the class of models considered in their paper does not include the model developed here, their remarks can readily be adapted.

7. Summary

At the core of the model described here is the concept that an individual's preferences are driven by his or her latent state, which we idealize as a strict weak order. The empty set, a special kind of strict weak order, is the 'neutral state', and serves as a reference: likes and dislikes are evaluated with respect to the neutral state. Over time, the state of an individual may change under the influence of quanta of information delivered by the environment. A quantum is called a 'token.' The potential effect of a token is to transform one state into another, neighboring the first. Profound changes may gradually result from the accumulation of such quantum changes. The tokens were classified into four categories: 'positive', 'negative', and their 'opposites', depending on their impact on the current state. The possible effect of a positive or negative token is to drive an alternative to the top or to the bottom of the strict weak order representing the state. The possible effect of the opposite is to drive the state toward the neutral state.

This qualitative version of the model was specialized, for the special case of three alternatives, by axioms casting the model as a (continuous parameter) Markov process, the succession of states being a homogeneous random walk on the collection of states. We supposed that the times of occurence of the tokens is governed by a homogeneous Poisson process. The effect of a token τ on a state $\succ$ was formalized by an operation $\succ \circ \tau = \succ'$ defined by Figure 1 (see also Figure 3). Quantitative results were stated in the form of four Theorems, which together yield exact predictions for the frequencies of pairs $(\succ, \succ')$ of strict weak orders obtained from a sample of respondents in two successive polls taken at times t and $t + \delta$, for large t (see Equation 18). Similar predictions could easily be derived for the general case of n polls taken at some arbitrarily chosen times $t_1 < t_2 < \cdots < t_n$, for large t_1.

Finally, we recall our earlier remark that the model presented here is but one of many possibilities falling under the same basic principles. Other related models are considered in Falmagne (1996) and Falmagne and Doignon (in press) (see also Doignon & Falmagne, in press). A general theory is presented in Falmagne (in press). We chose to present this particular model in view of its successful application to some important data, as reported in Regenwetter et al. (1995). As demonstrated in this paper, this type of model offers a powerful class of tools for the analysis of rich data. Further research along the lines of this paper would proceed first by analyzing a number of available panel data. The result of these analyses would dictate appropriate theoretical elaborations of these models such as those outlined in Section 5.

8. Proofs

Sketch of Proofs of Theorems 1-4. These proofs are similar to those of corresponding results in Falmagne and Doignon (in press). Accordingly, only an outline of the arguments is included here.

We suppose that the three Axioms [I], [T] and [L] hold. Theorem 1 is a straightforward consequence. To prove Theorem 2, notice that

$$\begin{aligned}
&\mathbb{P}(\boldsymbol{S}_{t+\delta} = \succ' \,|\boldsymbol{S}_t = \succ) \\
&= \sum_{k=0}^{\infty} \mathbb{P}(\boldsymbol{S}_{t+\delta} = \succ' \,|\boldsymbol{N}_{t,t+\delta} = k, \boldsymbol{S}_t = \succ)\, \mathbb{P}(\boldsymbol{N}_{t,t+\delta} = k|\boldsymbol{S}_t = \succ) \\
&= \sum_{k=0}^{\infty} \mathbb{P}(\boldsymbol{S}^*_{n+k} = \succ' \,|\boldsymbol{S}^*_n = \succ)\, \mathbb{P}(\boldsymbol{N}_{t,t+\delta} = k)\,.
\end{aligned}$$

With $p^*_{\succ,\succ'}(k)$ defined by Equation 9 and using Equation 5, Theorem 2 obtains.

The proof of Theorem 3 follows standard lines. The Markov chain (S^*_n) is irreducible and aperiodic, and thus has a unique stationary distribution. To prove that this stationary distribution is that given by Equation 12, we use the well known fact that if $(m_{R,S})_{R,S\in\mathcal{S}}$ is the transition matrix of a regular Markov chain on a finite set $\mathcal{S}$, and $\pi : R \mapsto \pi_R$ is a probability distribution on $\mathcal{S}$ satisfying the condition:

$$\forall R, S \in \mathcal{S}, \quad \pi_R \cdot m_{R,S} = \pi_S \cdot m_{S,R},$$

then π is the unique stationary distribution of the Markov chain. Writing K for the denominator of Equation 12 we only have to show that, for all $\succ, \succ' \in \mathcal{S}$ and distinct $i, j, k \in \{1, 2, 3\}$, we have

$$\frac{1}{K}\,\xi_i(\succ)\xi_j(\succ)\xi_k(\succ)p^*_{\succ,\succ'} = \frac{1}{K}\,\xi_i(\succ')\xi_j(\succ')\xi_k(\succ')p^*_{\succ',\succ}. \tag{19}$$

We check the case

$$\succ = [i \succ j, i \succ k], \qquad \succ' = [i \succ' j \succ' k]$$

and leave the others to the reader. Cancelling the denominators, Equation 19 specializes into

$$\tau_i\tilde{\beta}_i\tilde{\tau}_j\tilde{\beta}_j\tilde{\tau}_k\tilde{\beta}_k \cdot \beta_k = \tau_i\tilde{\beta}_i\tilde{\tau}_j\tilde{\beta}_j\tilde{\tau}_k\beta_k \cdot \tilde{\beta}_k,$$

which is trivial. The proof of Theorem 4 is immediate.

The Reparametrization in Remark 2 after Theorem 3. Dividing the numerator and the denominator of Equations 14-17 by $\tilde{\tau}_i\tilde{\beta}_i\tilde{\tau}_j\tilde{\beta}_j\tilde{\tau}_k\tilde{\beta}_k$ and defining

$$T_i(\succ) = \begin{cases} \frac{\tau_i}{\tilde{\tau}_i} & \text{if } \xi_i(\succ) = \tau_i\tilde{\beta}_i \\ 1 & \text{otherwise,} \end{cases}$$

$$B_i(\succ) = \begin{cases} \frac{\beta_i}{\tilde{\beta}_i} & \text{if } \xi_i(\succ) = \tilde{\tau}_i\beta_i \\ 1 & \text{otherwise,} \end{cases}$$

Equation 12 becomes

$$p_\succ = \frac{\Pi_{i=1}^3 T_i(\succ)B_i(\succ)}{\sum_{\succ'\in\mathcal{S}}\Pi_{i=1}^3 T_i(\succ)B_i(\succ)}.$$

More explicitly, writing C for the denominator in the above equation, the four generic expressions, Equations 14-17, become

$$\begin{aligned} p_{[i\succ j\succ k]} &= \frac{1}{C}\,T_i(\succ)B_k(\succ), \\ p_{[i\succ j, i\succ k]} &= \frac{1}{C}\,T_i(\succ), \\ p_{[j\succ i, k\succ i]} &= \frac{1}{C}\,B_i(\succ) \\ p_{[\emptyset]} &= \frac{1}{C}. \end{aligned}$$

The right members in these four expressions can be computed from the 6 parameters $T_i(\succ)$, $B_i(\succ)$, $i = 1, 2, 3$, which must be estimated from the data. Note, however, that when the data of two or more polls is considered, then the full quota of parameters is involved in the computation of the transitions probabilities $p^*_{\succ,\succ'}$ between states (c.f. Theorem 4).