

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Multilevel Hidden Markov Models for Behavioral Data: A Hawk-and-Dove Experiment

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Maruotti, A., Fabbri, M., Rizzolli, M. (2022). Multilevel Hidden Markov Models for Behavioral Data: A Hawk-and-Dove Experiment. MULTIVARIATE BEHAVIORAL RESEARCH, 57(5), 825-839 [10.1080/00273171.2021.1912583].

Availability:

This version is available at: <https://hdl.handle.net/11585/900936> since: 2022-11-09

Published:

DOI: <http://doi.org/10.1080/00273171.2021.1912583>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Maruotti, A., Fabbri, M., & Rizzolli, M. (2022). Multilevel hidden Markov models for behavioral data: A Hawk-and-Dove experiment. *Multivariate Behavioral Research*, 57(5), 825-839.

The final published version is available online at:

<https://doi.org/10.1080/00273171.2021.1912583>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

MULTILEVEL HIDDEN MARKOV MODELS FOR BEHAVIORAL DATA: A HAWK-AND-DOVE EXPERIMENT

BY ANTONELLO MARUOTTI^{*†}, MARCO FABBRI[‡] AND MATTEO RIZZOLLI[†]

University of Bergen^{*}, *Libera Università Maria Ss. Assunta*[†], *University of
Amsterdam*[‡]

XXX

1. Introduction. This paper is motivated by the statistical analysis of binary longitudinal data from an experimental study defined through the Hawk-and-Dove game initially introduced by [Maynard Smith and Parker \(1976\)](#) to investigate animals' behavior, and more recently extended to humans ([Sugden, 1989, 2004](#); [Waldron, 2013](#); [Rose, 2014](#)) in the game-theory literature. In the standard Hawk-and-Dove game two subjects face a *symmetric* situation in which they can choose between two strategies: to play hawk, i.e fighting aggressively for an asset; to play dove, i.e retreating the fight, if faced with major escalation, or sharing the asset, if not faced with escalation.

We investigate the evolution of an aggressive, fighting behavior, as playing hawk, and test the role of possession, property and further behavioral characteristics, with respect to such a behavior by introducing an *asymmetric* situation. Our treatment, possession-related, variation concerns the way the initial claim to the asset is established. We manipulate the type of information provided and the process of acquisition, creating an asymmetry in the game play. We aim at showing the variation in the probability of hawkish behaviors whenever the information is based on possession but also meritorious or, without merit or possession, for the presence of a *bullying* behavior.

To our knowledge, we are the first to test whether the hawkish behavior emerges in a lab setting once the possessor-intruder asymmetry is introduced. By having both possessory and non-possessory treatments, we are able to investigate that possession may be indeed a superior coordinating asymmetry. Moreover, meritoriousness may also play an important part in fostering the subjects behavior. At last, the evolution of the hawk behav-

MSC 2010 subject classifications: Primary 62-07, 62H25; secondary 62P12

Keywords and phrases: Hidden Markov models, Initial conditions, Conditional models, Possession

ior over time is investigated by looking at how the subjects learn to play from one round to another. With this design, we can address the following research questions

- Does owning/property induce more aggressive behaviors?
- Is merit important in establishing subjects behavior?
- Does possession trigger the establishment of property, and this favors a hawkish behavior?
- Is there a dependence over time in the subjects behavior?

In our experiment, subjects interact in groups of six over **ten** rounds. In each round, three pairs are randomly matched to interact under the Hawk-and-Dove setting. The matching is unknown to the players. Hence, outcomes, i.e. to play hawk or dove, are measured as binary longitudinal hierarchical data: repeated measurements are collected for subjects clustered into different groups. Thus, an *ad-hoc* statistical modeling should be considered to properly account for all data features to provide adequate answers to research questions defined above. As a general concern, when analyzing longitudinal hierarchical data, components which need to be described by a model the dependence of the variables on covariates, serial dependence, heterogeneity in the individuals/units and at the different levels of the hierarchy. Of course, the estimation of covariate effects on a response variable is often of major interest, while longitudinal correlations are typically viewed as nuisance parameters. However, the association structure could be of interest by itself, as we may be interested in understanding the nature of the stochastic dependence among the measured outcomes. In this respect, we introduce an adequate definition of the association structure distinguishing between *true* and *apparent* contagion, also known as *state dependence* and *heterogeneity*, respectively. In the former case, the occurrence of an event changes the probability of the subsequent occurrence of similar events: in a longitudinal setting, actual and future outcomes are directly influenced by past values, which cause a substantial change over time in the corresponding distribution. The latter case arises when individuals are drawn from heterogeneous populations, each population having a different, propensity to a certain event. In practice, true contagion is modeled by simply including the lagged outcome as an additional covariate, and the apparent contagion is accounted for by including, possibly time-varying, random effects in the model specification (Aitkin and Alfó, 1998; Alfó and Aitkin, 2000; Aitkin and Alfó, 2003; Fotouhi, 2005; Skrondal and Rabe-Hesketh, 2014; Bianconcini and Bollen, 2018). Unfortunately, the corresponding maximum likelihood estimators can be quite severely inconsistent owing to the initial conditions

problem. The initial response at the start of the observation period is affected by the random intercept and presample responses, and ignoring this endogeneity leads to inconsistent estimation (Heckman, 1981).

This paper presents a model for the analysis of longitudinal hierarchical binary outcomes which extends standard random effects models to simultaneously account for serial dependence and heterogeneity at the different levels of the hierarchy. In detail, we develop a multilevel approach in which the unobserved heterogeneity at the different level of the hierarchy, i.e. unobserved subjects' and groups' behaviors, is assumed discrete and modeled through the inclusion of random effects (Aitkin, 1999; Bartolucci and Farcomeni, 2009). This involves an additional step in the model selection, because the support of these distributions is not known in advance; it must be selected by evaluating the goodness of fit that different supports obtain. There are, however, several advantages that compensate such complication. First, the random-effect model reduces to a finite mixture model with a computationally tractable likelihood function. Second, the possibly inappropriate and unverifiable parametric assumptions about the distribution of the random effects are avoided. Third, the outcomes are clustered in a finite number of latent classes that can be interpreted as typical behaviors. A conditional (to the initial conditions) model is introduced and proper inference is conducted. A shape change in the random effects distribution is considered by directly modeling the effect of the initial conditions on the evolution of behaviors. The observed measurements are modeled through a generalized hierarchical linear mixed model. For the maximum likelihood estimation of the proposed model, we use an EM-based algorithm deriving, and slightly modifying, recursions from the Baum-Welch algorithm (Baum et al., 1970), widely adopted in the hidden Markov models literature (Bartolucci et al., 2013). However, different methods can be used to provide parameter estimates (Bulla and Berzel, 2008).

Several models are fitted and compared to properly infer the observed and latent structures driving the data. Starting from a simple model without covariates to more complex multilevel models, allowing for true and spurious contagion. The rest of the paper is as follow. Section 2 describes the Hawk-and-Dove experiment, along with the considered *treatments*. An initial description and assessment of the data is provided to guide the reader into the collected data. In Section 3, we introduce a novel methodology, that can be cast in the literature on multilevel modelling and mixed hidden Markov models. The initial condition problem is also briefly sketched and solutions proposed. Section 4 is devoted to likelihood inference, and some operational aspects are further discussed.

		Subject 2	
		Hawk	Dove
Subject 1	Hawk	-25	50
	Dove	0	15

TABLE 1

The Hawk-and-Dove game: payoffs

2. The Hawk-and-Dove experiment.

2.1. *Data collection and treatment definition.* In our experiment, subjects interact in groups of six over **ten** rounds. In each round, three pairs are randomly matched to interact in two distinct activities. In the first part of the round, asymmetry is introduced, while in the second part the two subjects of each pair play the Hawk-and-Dove game and contend an amount of experimental monetary units (see Table 1 for the resulting payoffs). This second part is always the same across rounds and treatments.

It is in the first part that our manipulation takes place. Our treatment variation concerns the way the initial claim is established, that is to say, the nature of the asymmetry we provide. We have a total of five treatments organized along two main dimensions. The first dimension concerns the type of information provided (the asymmetry is either *possessory* or *colored*), and the other dimension concerns the process of acquisition (the asymmetry is either *arbitrary* or *meritorious*). In designing our possessory treatments, we have mimicked some well-known mechanisms of property acquisition that are also almost universally enforced in property law. Among all the possible ways property can be acquired, we have focused on reproducing three benchmarks: *Gift*, *Treasure Trove*, and *Labor*.

In the ***Gift*** treatment, the computer randomly assigns 50 tokens to one of the two subjects (as a manna from heaven), who becomes the *possessor*. The other subject (the *intruder*) does not receive any tokens. The asymmetry is therefore possessory and arbitrary.

In the ***Labor*** treatment, participants have to perform an individual effort task following Gill and Prowse (2012). Each participant has one minute to move the cursors of as many 0–100 scale sliders as possible to the position indicated by the computer. For each matched pair of subjects, the one that correctly positions the highest number of cursors gains possession of 50 tokens. The asymmetry is therefore possessory and meritorious. Because this activity does not require any particular ability or knowledge to complete, it is likely that participants perceive the endowment gained through individual performance to be correlated with individual effort. By using a meritorious mechanism of acquisition based on effort, this treatment inten-

tionally mimics labour as an almost universal mechanism for legitimizing property (Locke, 1980; Henry, 1999).

In the *Treasure Trove* treatment, at the beginning of each round subjects participate in a treasure trove contest. Each pair sees the same 25 squares on the computer screen and can uncover their content by pressing on each of them. Hidden behind one of the squares is a 50-token treasure in the form of a code composed of numbers and letters. Whoever registers its trove in the dedicated filling area at the bottom of the screen takes possession of the 50-token treasure that will then be contested in the second part of each round. The asymmetry is therefore possessory and meritorious.

For the non-possessory treatments, we have followed some previous designs that based the asymmetry on an assigned *color*: one subject enters the Hawk-and-Dove game being red and the other enters being blue. We thus have an arbitrary colored treatment (*Lucky Red*) and a meritorious colored treatment (*Master Red*). In the *Lucky Red* treatment, each subject is randomly labeled either *Red* or *Blue* at the beginning of each round so that a asymmetry is introduced along the lines of Hargreaves-Heap and Varoufakis (2002). However, this color label is not related to the initial possession of the 50 tokens; in fact nobody possesses the tokens before the contest.

In the *Master Red* treatment, subjects participate in a treasure hunt at the beginning of each round as in the *Treasure Trove* treatment. Whoever registers their trove in the dedicated filling area at the bottom of the screen is assigned color Red (Blue). Notice that the assignment of this color label is not related to the initial possession of the 50 tokens; in fact, nobody possesses the tokens before the contest.

2.2. Data description. We conducted laboratory sessions with 12 or 18 subjects each, for a total of 474 participants clustered into 79 groups. Each subject participated one session only. The experiment was conducted using computer interfaces, and, to program the experiment, we used the software Z-tree (Fischbacher, 2007). The vast majority of participants were graduate and undergraduate students at the University and were recruited using the online system ORSEE (Greiner, 2015). At the beginning of each session, instructions were read aloud by the experimenter to ensure common knowledge. Before the experiment started, all the participants had to correctly answer some control questions. Throughout the reading of the instructions and the control questions stage, participants had the opportunity to ask the experimenter questions in private. After the experiment ended, each subject was asked to fill in a questionnaire reporting socio-demographic characteristics and measuring individual risk preferences, logical abilities, and the level

of impulsivity. Communication among participants was not allowed during the experiment.

A preliminary assessment of the behavior of the groups in being hawkish may be depicted by looking at Figure 1. We collapse the number of hawks into three categories to simplify data visualization. At the beginning of the experiment, only players into few groups are clearly in favour of a hawkish behavior, as almost all are playing hawk. In the majority of groups, players are equally divided by those who play hawk and those who play dove. However, after few periods, the situation is more complex. On the one side, it is possible to identify groups which find their *equilibrium*, i.e. the number of hawk and dove players remain constant over time; on the other hand, heterogeneity in groups' behaviors can be inferred, with a small increase in hawkish behaviors over time, and clear paths cannot be identified. However, something can be said about the dynamics of groups' behaviors, by looking at the empirical transition matrix, i.e. at the frequencies of moving between the three categories [0 – 2], [3 – 4], [5 – 6]. Persistence is not the norm for two of the three categories. If a small number of players ([0-2]) played hawk at time t , at time $t + 1$ such a number would likely increase to [3-4] with *probability* 0.61 and, rarely, to [5-6] (only in the 8% of the cases). Similarly, if a large number of subjects played hawk at time t , it is likely that this number decreases at time $t + 1$ passing to [3-4] or, rarely, to [0-2]. The mid category is the persistent one with the 61% of cases that the same number of hawkish is observed at two consecutive times. What is also rather clear is that if a group is incline to play hawk, it does not transit toward a dove behavior between two consecutive times, but passes from being very aggressive to be less aggressive first, i.e. from category [5-6] to [3-4], the viceversa happens for groups more likely to play dove.

$$\begin{bmatrix} 0.31 & 0.61 & 0.08 \\ 0.19 & 0.61 & 0.20 \\ 0.09 & 0.65 & 0.26 \end{bmatrix}$$

Hawkish behaviors may depend on several factors and, by a preliminary assessment, slightly increase over time. This is also depicted by looking at Figure 2. A small increase in the overall hawkish behavior is observed, but significant differences are observed between owner and non-owner subjects. Subjects owning the experimental tokens are more likely to defend them and, accordingly, they are more likely to play hawk. More than 60% of the owners tend to play hawk, with peaks close to the 70%. Interestingly, for non-owner subjects the percentage of subject playing hawk is also rather high, and above 50%, suggesting that, on average, the hawk behavior is

FIG 1. *Number of subjects playing hawk over time per group*


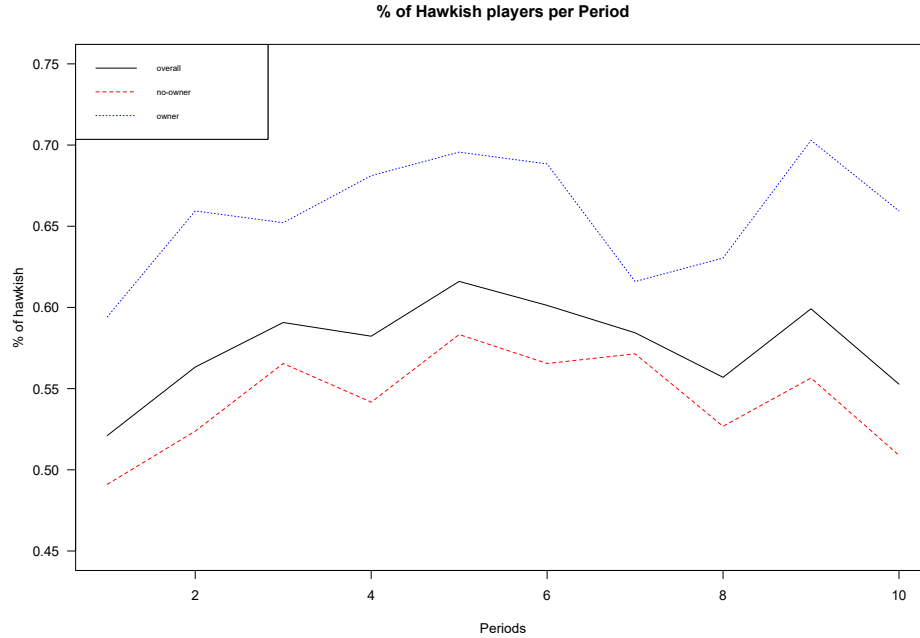


FIG 2. *Percentage of subjects playing hawk over time, owner vs. non-owner subjects*

preferred to the dove one.


We further expect that the asymmetry introduced by the treatments would also play a role in subjects' behavior. Figure 3 shows the percentage of hawks by treatment. Differences can be observed among treatments, though the observed variability, quite high at the beginning of the experiment, tends to decrease over time, as players learned to play, and get info from previous rounds.

3. Modelling.

3.1. The multilevel hidden Markov model. The data considered in this paper are in the form of panels of subjects behaviors, assigned/clustered within playing groups. A clear hierarchy can be easily detected in the data structure and different sources of heterogeneity should be tackled at the different levels of the hierarchy.

Formally, the response variable, i.e. behave hawkish, Y_{ntg} assumes two values only, $\{0,1\}$, and is observed on N ($n = 1, 2, \dots, N_g$) individuals clustered into G ($g = 1, 2, \dots, G$) groups at T_n ($t = 1, 2, \dots, T_n$) occasions. Since

FIG 3. *Percentage of subjects playing hawk over time by treatment*



repeated measurements belonging to the same subject assigned to a specific group are likely to be correlated, a modelling framework accounting such dependence structure should be specified, for valid inferences. A natural way to account for correlated measurements is via multilevel random effects models, as they provide a flexible way to deal with complex data structures. We assume that, conditionally on a set of available exogenous subject-specific covariates $\mathbf{x}_{ntg} = \{x_{ntg1}, \dots, x_{ntgP}\}$, the responses are generated by combining individual- and group-specific random components. The individual-specific random component is specified by an array $\mathbf{b} = \{\mathbf{b}_1, \dots, \mathbf{b}_N\}$ of N time-dependent trajectories, with $\mathbf{b}_n = \{b_{n1}, \dots, b_{nT_n}\}$, following a distribution $p(\mathbf{b})$. The group-specific random component is, instead, specified by a time-constant random effect v_g , independent of \mathbf{b}_n , following a certain distribution $p(\mathbf{v}) = p(v_1, \dots, v_G)$.

We further assume that the responses are conditionally independent given the individual- and the group-specific random effects, or, in other words, that the conditional distribution of all the responses y_{ntg} is a product of univariate conditional distributions, say

$$p(\mathbf{y} \mid \mathbf{b}, \mathbf{v}) = \prod_{g=1}^G \prod_{n=1}^{N_g} \prod_{t=1}^{T_n} p(y_{ntg} \mid \mathbf{b}, \mathbf{v}).$$

Under these assumptions, we have that

$$(3.1) \quad p(\mathbf{y}, \mathbf{b}, \mathbf{v}) = \prod_{g=1}^G p(v_g) \prod_{n=1}^{N_g} p(\mathbf{b}_n) \prod_{t=1}^{T_n} p(y_{ntg} \mid \mathbf{b}_n, v_g)$$

is fully specified by defining

- (a) the conditional distribution of the response variable given covariates and random effects,
- (b) the distribution of the group-specific random effects and
- (c) the distribution of the individual-specific trajectories of random effects.

We assume that

$$Y_{ntg} \sim \text{Bernoulli}(\lambda_{ntg})$$

where

$$(3.2) \quad \log \left(\frac{\lambda_{ntg}}{1 - \lambda_{ntg}} \right) = \mathbf{x}'_{ntg} \boldsymbol{\beta} + \alpha y_{n(t-1)g} + v_g + b_{nt}$$

where $\boldsymbol{\beta}$ is a vector of regression coefficients and α is the regression coefficient associated to the lagged response, introduced to account for true contagion.

To specify the distributions of the random terms, we assume that the random effects are drawn from discrete distributions with a finite number of mass points. At the group-level, we assume that the random effects v_g are independently drawn for each group from a discrete distribution, say

$$(3.3) \quad v_g \sim p_M(\mathbf{v}; \boldsymbol{\pi})$$

where $p_M(\mathbf{v}; \boldsymbol{\pi})$ depends on a vector of M support points $\mathbf{v} = (v_1, \dots, v_M)$ with mass probabilities $\boldsymbol{\pi} = (\pi_1, \dots, \pi_M)$. We further assume that the individual-specific process follows a first-order finite-state Markov chain with state-space $B = (b_1, \dots, b_K)$. The chain is fully known up to an initial distribution $\boldsymbol{\delta} = (\delta_1, \dots, \delta_K)$, $\delta_k = \Pr(b_{n1} = b_k)$, and a $K \times K$ transition probability matrix \mathbf{Q} whose elements are given by $q_{nthk} = \Pr(b_{nt} = b_k \mid b_{n(t-1)} = b_h)$, $t > 1; h, k = 1, 2, \dots, K; \sum_{h=1}^K q_{nthk} = 1$. In its basic specification, we will consider a homogenous transition probability matrix, i.e. $q_{nthk} = q_{hk}$. Accordingly, for each individual, the sequence \mathbf{b}_n is drawn from a discrete distribution, say

$$(3.4) \quad \mathbf{b}_n \sim p_K(\mathbf{b}; \boldsymbol{\delta}, \mathbf{Q}).$$

Bearing in mind that this model can be also cast in the literature on the mixed hidden Markov models (Maruotti, 2011) and is related to the approaches introduced by Bartolucci et al. (2014); Bartolucci and Lupparelli (2016); Leos-Barajas et al. (2017); DeRuiter et al. (2017); Fr urwirth-Schnatter et al. (2018); Montanari et al. (2018); Adam et al. (2019); Zhang and Chang (2019) and Lagona et al. (2015), this model includes some popular approaches to longitudinal data analysis as particular cases. For example, when the number of the hidden states K is equal to 1, the model reduces to a finite mixture model, in which data are clustered within G classes. On the other side, when the distribution $p(\mathbf{v})$ concentrates the whole probability mass on the origin, $p(\mathbf{v} = 0) = 1$, the model reduces to a hidden Markov model.

3.2. The initial conditions problem and the conditional approach. The joint distribution of the observed responses for subject n , given the group-specific random effect v_g is given by

$$(3.5) \quad p(y_{n1g}, \dots, y_{nT_n g} \mid v_g) = \sum_{\mathbf{b}_n} p_K(\mathbf{b}_n; \boldsymbol{\delta}, \mathbf{Q}) p(y_{n1g} \mid b_{n1}, v_g) \prod_{t=2}^{T_n} p(y_{ntg} \mid y_{n(t-1)g}, b_{nt}, v_g)$$

where $\sum_{\mathbf{b}_n}$ extended to all the possible configurations of \mathbf{b}_n .

Nevertheless, equation (3.5) is not determined from the model assumptions since the distribution $p(y_{n1g} | b_{n1}, v_g)$ is not specified from the model and the full likelihood is not available. Inference can be highly sensitive to misspecification of $p(y_{n1g} | b_{n1}, v_g)$. An alternative is to estimate regression parameters by maximizing the likelihood conditional on the first outcome Y_{n1g} , i.e. considering the distribution of $p(y_{n2g}, \dots, y_{nTg} | y_{n1g})$:

$$(3.6) \quad p(y_{n2g}, \dots, y_{nTg} | y_{n1g}, v_g) = \sum_{\mathbf{b}} p_K(\mathbf{b}_n, \boldsymbol{\delta}, \mathbf{Q}) \prod_{t=2}^{T_n} p(y_{ntg} | y_{n(t-1)g}, b_{nt}, v_g).$$

Bearing in mind that \mathbf{b}_n is a subject-specific effect shared by all subject n 's outcomes, and as well by Y_{n1g} , \mathbf{b}_n and Y_{n1g} cannot be assumed independent. If independence is assumed, the resulting estimator is inconsistent because the initial response at time $t = 1$ is treated as exogenous, while it is clearly endogenous, giving rise to the initial conditions problem.

The basic idea is to re-express the conditional model we are dealing with by allowing for the dependence between Y_{n1g} and \mathbf{b}_n . The underlying hypothesis is that the influence of Y_{n1g} on \mathbf{b}_n can be fully modelled as a change in the location and shape of \mathbf{b}_n . Following Wooldridge (2005) and Aitkin and Alfó (2003), let us assume that

$$(3.7) \quad \mathbf{b}_n = \mathbf{b}_n^* + E[\mathbf{b}_n | y_{n1g}]$$

where

$$E[\mathbf{b}_n | Y_{n1g}] = \tilde{\alpha} y_{n1g}.$$

Then, equation (3.2) can be rewritten as

$$(3.8) \quad \log \left(\frac{\lambda_{ntg}}{1 - \lambda_{ntg}} \right) = \mathbf{x}'_{ntg} \boldsymbol{\beta} + \alpha y_{n(t-1)g} + \tilde{\alpha} y_{n1g} + v_g + b_{nt}.$$

In the following, without loss of generality, we will drop the superscript * and refer to \mathbf{b}_n^* as \mathbf{b}_n .

Furthermore, we allow for a more general change in the shape of the subject-specific random effects \mathbf{b}_n , i.e. we assume that Y_{n1g} affects the Markov chain parameters as well. Accordingly, we rewrite expression (3.6):

$$(3.9) \quad p(y_{n2g}, \dots, y_{nTg} | y_{n1g}, v_g) = \sum_{\mathbf{b}} p_K(\mathbf{b}_n | y_{n1g}, \boldsymbol{\delta}_{y_{n1g}}, \mathbf{Q}_{y_{n1g}}) \prod_{t=2}^T p(y_{ntg} | y_{n(t-1)g}, y_{n1g}, b_{nt}, v_g).$$

where the conditional distribution $p_K(\mathbf{b}_n \mid y_{n1g}; \boldsymbol{\delta}_{y_{n1g}}, \mathbf{Q}_{y_{n1g}})$ is different from $p_K(\mathbf{b}_n; \boldsymbol{\delta}, \mathbf{Q})$. In detail, such a dependence on Y_{n1g} leads to

$$(3.10) \quad \delta_{nk} = \Pr(b_{n2} = b_k \mid y_{n1g}) = \frac{\exp(\gamma_{0k} + \gamma_{1k}y_{n1g})}{\sum_{h=1}^K \exp(\gamma_{0h} + \gamma_{1h}y_{n1g})}$$

and similarly we link the initial outcome and the entries of the transition probability matrix

$$(3.11) \quad q_{nhk} = \Pr(b_{nt} = b_k \mid b_{n(t-1)} = b_h, y_{n1g}) = \frac{\exp(\phi_{0hk} + \phi_{1hk}y_{n1g})}{\sum_{k=1}^K \exp(\phi_{0hk} + \phi_{1hk}y_{n1g})}$$

Subjects do not share the same latent structure, except if they have the same value of y_{n1g} . Indeed, different homogeneous (over time) Markov chains have been defined conditionally on y_{n1g} .

4. Maximum likelihood estimation. The multilevel logistic model illustrated in Section 3.1 depends on a vector $\boldsymbol{\theta}$ of parameters that includes four components:

- (a) the fixed effects $\boldsymbol{\beta}$ and $\boldsymbol{\alpha} = \{\alpha, \tilde{\alpha}\}$;
- (b) the support points (v_1, \dots, v_m) and the related probabilities $\boldsymbol{\pi}_M$ of the random effects distribution at the group level;
- (c) the support points (b_1, \dots, b_K) of the time-varying random effects;
- (d) the fixed effects $\boldsymbol{\gamma}$ and $\boldsymbol{\phi}$ related to the initial probabilities collected in $\boldsymbol{\delta}_{y_{n1g}}$ and the transition probabilities collected in $\mathbf{Q}_{y_{n1g}}$, respectively.

The maximum likelihood estimate of $\boldsymbol{\theta}$ is the maximum point of the conditional likelihood function

$$(4.1) \quad L(\boldsymbol{\theta} \mid \mathbf{y}_1) = \prod_{g=1}^G \sum_{\mathbf{v}} p_M(\mathbf{v}; \boldsymbol{\pi}) \prod_{n=1}^{N_g} \sum_{\mathbf{b}} p_K(\mathbf{b}; \boldsymbol{\delta}_{y_{n1g}}, \mathbf{Q}_{y_{n1g}}) \prod_{t=2}^{T_n} p(y_{ntg} \mid y_{n(t-1)g}, y_{n1g}, b_{nt}, v_g; \boldsymbol{\beta}, \boldsymbol{\alpha})$$

Expression (4.1) can be efficiently computed by an extension of the forward recursion, which is very well-known in the HMM literature (Welch, 2003). To maximize the likelihood, we implement a version of the expectation-maximization (EM) algorithm, which is facilitated by the independence assumption between the time-varying subject-specific and the group-specific random effects.

The algorithm is based on the definition of the so-called complete-data log-likelihood function, obtained by considering the sampling distribution of

both the observed and the unobserved quantities, i.e. the unknown states and groups memberships. Treating these quantities as missing values, reflecting different sources of incomplete information, we define the complete-data log-likelihood function as

$$\ell_c(\boldsymbol{\theta} \mid \mathbf{y}_1) =$$

$$(4.2) \quad = \sum_{g=1}^G \sum_{m=1}^M \eta_{gm} \log(\pi_m)$$

$$(4.3) \quad + \sum_{n=1}^N \sum_{k=1}^K \xi_{n2k} \log(\delta_{nk})$$

$$(4.4) \quad + \sum_{n=1}^N \sum_{k=1}^K \sum_{h=1}^K \sum_{t=3}^{T_n} \zeta_{nthk} \log(q_{nhk})$$

$$(4.5) \quad + \sum_{g=1}^G \sum_{m=1}^M \sum_{k=1}^K \sum_{n=1}^N \sum_{t=2}^{T_n} \eta_{gm} \xi_{ntk} \log(p(y_{ntg} \mid y_{n(t-1)g}, y_{n1g} b_{nt}, v_g; \boldsymbol{\beta}, \boldsymbol{\alpha}))$$

where the variable $\eta_{gm} = I(v_g = v_m)$ is an indicator variable equal to 1 if group g is clustered in cluster m , $\zeta_{nthk} = I(b_{nt} = b_k, b_{n(t-1)} = b_h)$ is an indicator variable equal to 1 if subject n belongs to state h at time $t-1$ and to state k at time t , and $\xi_{ntk} = I(b_{nt} = b_k)$ equals 1 if subject n at time t belongs to state k and 0 otherwise.

In the E-step, the conditional expected value of terms (4.2)–(4.5) is simply computed by a plug-in of the expected values of η_{gm} , ξ_{ntk} and ζ_{nthk} given the observed data and the current value of the parameters. Such quantities can be computed by means of an appropriate forward-backward recursion adapted from the mixed hidden Markov model literature (Maruotti, 2011). In the M-step, the conditional expected values of terms (4.2)–(4.5) are maximized separately. In particular, at iteration $r+1$, the maximum with respect to π_m has a closed form solution

$$\pi_m^{(r+1)} = \frac{\sum_{g=1}^G \mathbb{E}(\eta_{gm} \mid \mathbf{y}, \boldsymbol{\theta}^{(r)})}{G}.$$

The maximum with respect to δ_{nk} and q_{nhk} is obtained as solutions of the following M-step equations, respectively,

$$(4.6) \quad \sum_{n=1}^N \sum_{k=1}^K \mathbb{E}(\xi_{n2k} \mid \mathbf{y}, \boldsymbol{\theta}^{(r)}) \frac{\partial \log(\delta_{nk})}{\partial \gamma_k} = 0$$

and

$$(4.7) \quad \sum_{n=1}^N \sum_{t=3}^{T_n} \sum_{k=1}^K \mathbb{E}(\zeta_{nthk} \mid \mathbf{y}, \boldsymbol{\theta}^{(r)}) \frac{\partial \log(q_{nhk})}{\partial \phi_h} = 0$$

which are weighted sums of K equations with weights $\mathbb{E}(\xi_{n2k} \mid \mathbf{y}, \boldsymbol{\theta}^{(r)})$ and $\mathbb{E}(\zeta_{nthk} \mid \mathbf{y}, \boldsymbol{\theta}^{(r)})$, see also [Maruotti and Rocci \(2012\)](#). Similarly, denoting with $\boldsymbol{\theta}^* = \boldsymbol{\beta}, \alpha, \mathbf{b}, \mathbf{v}$ the parameters of the state-dependent distribution, we obtain the updated estimates of $\boldsymbol{\theta}^*$, solving the following equation

$$\sum_{g=1}^G \sum_{m=1}^M \sum_{k=1}^K \sum_{n=1}^{N_g} \sum_{t=1}^{T_g} \mathbb{E}(\eta_{gm} \mid \mathbf{y}, \boldsymbol{\theta}^{(r)}) \mathbb{E}(\xi_{ntk} \mid \mathbf{y}, \boldsymbol{\theta}^{(r)}) \log p(y_{ntg} \mid y_{n(t-1)g}, y_{n1g}, b_{nt}, v_g; \boldsymbol{\beta}, \boldsymbol{\alpha}) = 0.$$

We alternate the E and M steps repeatedly until the increment in the likelihood is less than a fixed, small amount.

4.1. Initialization of the algorithm, computational strategies and model selection. The algorithm may be trapped at a local maximum and, consequently, may fail to reach global maximum (because the maximization problem is likely to be non-convex). One simple way to alleviate the problem is to run the EM algorithm from multiple random starting points for a number of steps, then pick the one with the highest likelihood, and continue the EM from the picked point until convergence. Furthermore, the EM algorithm outlined above does not produce standard errors of the estimates, because approximations based on observed information matrix often requires a very large sample size. Thus, to obtain standard errors, we consider a parametric bootstrap approach, refitting 200 bootstrap samples simulated from the estimated model parameters. The approximate standard error of each model parameter is then computed.

In the proposed multilevel framework, model selection is essentially concerned with the choice of the order, that is, the number of hidden states K and mixture components M . This decision entails a preliminary exploration of a range of different values of K and M and a final choice requires typically a compromise between several different factors including model fitting, formal selection criteria, computational complexity, and the overall interpretability of results [Pohle et al. \(2017\)](#). In this application, we fit models with $K = \{2, 3, 4\}$ and $M = 2$, as we have noticed overfitting for greater values of M . Classical information criteria based on log-likelihood penalizations are adopted for order selection, like the AIC and BIC criteria.

To avoid the multinomial regressions defined by (4.6) and (4.7), we could exploit the binary nature of the initial conditions. Let $N_0 = \{n : Y_{n1g} =$

$0; n = 1, \dots, N\}$ and $N_1 = \{n : Y_{n1g} = 1; n = 1, \dots, N\}$ be the subsets of the N individuals having respectively $Y_{n1g} = 0$ and $Y_{n1g} = 1$, with $|N_0| + |N_1| = n_0 + n_1 = N$.

The likelihood then can be written as

$$L(\boldsymbol{\theta} \mid \mathbf{y}_1) = \prod_{g=1}^G \sum_{\mathbf{v}} p_M(\mathbf{v}; \boldsymbol{\pi}) \left\{ \prod_{n \in N_0} \sum_{\mathbf{b}} p_K(\mathbf{b}; \boldsymbol{\delta}_{y_{n1g}=0}, \mathbf{Q}_{y_{n1g}=0}) \prod_{t=2}^{T_n} p(y_{ntg} \mid y_{n(t-1)g}, y_{n1g}, b_{nt}, v_g; \boldsymbol{\beta}, \boldsymbol{\alpha}) \prod_{n \in N_1} \sum_{\mathbf{b}} p_K(\mathbf{b}; \boldsymbol{\delta}_{y_{n1g}=1}, \mathbf{Q}_{y_{n1g}=1}) \prod_{t=2}^{T_n} p(y_{ntg} \mid y_{n(t-1)g}, y_{n1g}, b_{nt}, v_g; \boldsymbol{\beta}, \boldsymbol{\alpha}) \right\}.$$

Differentiating the previous equation with respect to hidden chain parameters under the constraints and equating to zero the corresponding derivatives, the M-step reduces to

$$\delta_{nk; y_{n1g}=i}^{(m+1)} = \sum_{n \in N_i} \frac{\xi_{n2k; y_{n1g}=i}}{n_i}; \quad i = 0; 1$$

and

$$q_{nthk, y_{n1g}=i}^{(m+1)} = \sum_{n \in N_i} \frac{\zeta_{nthk, y_{n1g}=i}}{\sum_k \zeta_{nthk, y_{n1g}=i}}; \quad i = 0; 1.$$

5. Results. In this Section, we illustrate the results obtained from the proposed approach to the dataset about subjects behavior in a hawk and dove game described in Section 2. We fitted, and compared, several models in terms of penalized likelihood criteria, namely:

- M1: a simple HMM, with no covariates;
- M2: an HMM with treatments, gender and age as covariates;
- M3: an HMM with the lagged outcome, treatments, gender and age as covariates;
- M4: an HMM with the lagged outcome, treatments, gender and age as covariates, and initial conditions included in the model as described in Section 3.2;
- M5: a multilevel HMM with the lagged outcome, treatments, gender and age as covariates, initial conditions included in the model as described in Section 3.2 and group-specific random effects.

All models are fitted for $K = \{2, 3, 4\}$, Table 2 displays the log-likelihoods with the number of parameters and the values attained by the AIC and BIC criteria. According to both criteria, the multilevel model is chosen, confirming the adequacy of the proposed approach for the data at hand.

TABLE 2
Results from fitting different hidden Markov models with different values of K . The maximum log-likelihood of each model is denoted by ℓ .

M1					
ℓ	K	M	# parameters	AIC	BIC
-2753.576	2		5	5517.153	5537.959
-2706.514	3		10	5433.641	5474.641
-2706.388	4		17	5447.776	5517.517
M2					
ℓ	K	M	# parameters	AIC	BIC
-2713.548	2		12	5451.096	5501.03
-2664.858	3		17	5363.716	5434.457
-2664.849	4		24	5377.698	5477.567
M3					
ℓ	K	M	# parameters	AIC	BIC
-2709.726	2		13	5445.453	5499.549
-2662.456	3		18	5362.912	5435.814
-2659.805	4		27	5373.610	5482.962
M4					
ℓ	K	M	# parameters	AIC	BIC
-2618.314	2		17	5270.627	5341.368
-2607.647	3		28	5271.295	5387.809
-2604.554	4		43	5295.108	5474.040
M5					
ℓ	K	M	# parameters	AIC	BIC
-2289.710	2	2	20	4619.420	4702.644
-2279.557	3	2	32	4623.114	4756.272
-2279.546	4	2	49	4657.092	4869.991

The estimates of the parameters under the selected models are reported in Table 3.

Regression parameters associated to the considered covariates are rather consistent across the models and allow to draw some conclusions. Firstly, the *Gift* and *Treasure Trove* treatments do not differ from **Labor**, that is taken as reference treatment. So, players behavior does not change over different possessory treatments. The **Lucky Red** and **Master Red** treatments, instead, increase the probability of playing hawk. This is a clear indication that introducing asymmetry, i.e. establishing the possess of the tokens,

TABLE 3
Regression estimates from fitting different hidden Markov models (standard errors in brackets).

Variable/Parameter	M1	M2	M3	M4	M5
Fixed effects					
Intercept (β_0)	0.604	0.322 (0.326)	0.297 (0.332)	-1.683	-1.912
Lucky Red (β_1)	-	0.346 (0.170)	0.356 (0.173)	0.357 (0.121)	0.393 (0.121)
Gift (β_2)	-	-0.204 (0.191)	-0.211 (0.192)	-0.205 (0.114)	-0.184 (0.114)
Treasure Trove (β_3)	-	0.037 (0.167)	0.044 (0.170)	0.020 (0.114)	0.015 (0.114)
Master Red (β_4)	-	0.329 (0.171)	0.339 (0.173)	0.341 (0.118)	0.356 (0.118)
Owner (β_5)	-	0.916(0.103)	0.928 (0.106)	0.863 (0.093)	0.856 (0.093)
Male (β_6)	-	0.039 (0.109)	0.035 (0.110)	-0.014 (0.070)	0.036 (0.070)
Age (β_7)	-	-0.007 (0.023)	-0.007 (0.024)	0.017 (0.016)	0.021 (0.016)
Lag-outcome (α)	-	-	-0.050 (0.100)	0.083 (0.072)	0.079 (0.041)
Initial conditions ($\bar{\alpha}$)	-	-	-	2.609 (0.131)	2.769 (0.141)
Subject-specific random effects					
$b_{nt} = b_1$	-2.390	-2.677	-2.755	-1.340	-1.259
$b_{nt} = b_2$	-0.272	-0.251	-0.256	1.232	1.390
$b_{nt} = b_3$	3.345	3.031	3.035	-	-
Group-specific random effects					
$v_g = v_1$	-	-	-	-	0.080
$v_g = v_2$	-	-	-	-	-0.182

clearly affects behaviors. In particular, being *Lucky Red* and *Master Red* both non-possessory treatments, we conclude that aggressive behaviors arise when none of the players possess the tokens. At the same time, owning the tokens is one of the main driving variables for playing hawk, i.e. property makes the difference in players' behavior. Accordingly, we notice that if a resource (i.e. the tokens) is freely available players are more incline to fight, i.e. playing hawk, but this is also true that if the property is clearly defined, the owner would likely defend it.

The probability of playing hawk is smaller if players are aware that the tokens are assigned to any of the players, as a form of respect between players. In some sense, possession-related treatments indicate a coordinating system. This is in line with the so-called bourgeois strategy (Maynard Smith and Parker, 1976): respect others' belongings and expect others to respect theirs. At the same time it possible to estimate a significant difference between merit and non-merit treatments. In the *Treasure Trove* treatment the probability of playing hawk is higher than in the *Gift* treatment, where the possession is a *mana from heaven* and no merit causes the possession. In practice, if a player does nothing to merit the possession, then he/she less likely would play hawk.

Looking at players' behavior over time, *true* contagion plays a minor role. Previous choices have only a small effect on current ones, in particular with respect to the role played by treatments and property variables. Nevertheless, the choice at the baseline has a great impact on all future choices. A

hawk behavior at $t = 1$ very likely leads to a hawk behavior at subsequent times, indicating persistence of behaviors, not due to the way subjects play over time, but rather on an initial attitude.

Heterogeneity arises at both levels of the hierarchy. In the following, we comment only on the multilevel model M5. For the multilevel model, M5, estimates of the hidden Markov chain and mixture parameters are reported in Table 4. Two states and two components are estimated at the subject and group levels, respectively. States represent two different subjects, more or less aggressive, behaviors, and mixture component different ways of coordinating inside the groups. In State 1 ($\mathbf{b}_n = 1$) subjects with a dove propensity are clustered, while State 2 ($\mathbf{b}_n = 2$) clusters *bullying* attitudes. The two states are well separated, and as expected they communicate rarely (see the transition matrices in Table 4). So the behaviors at the subject level are well identified and characterized by strong persistence over time. The group effects have lower impacts on the probability of playing hawk. It deserves to be mentioned that around 31% of the groups (i.e. those clustered in the second mixture component) show a propensity of sharing, further reducing the probability of playing hawk and converge to more likely peaceful behavior. For the other groups, a small, rather marginal, effect is estimated; and this does not deserve to be further discussed.

At last, it is rather interesting to notice that the initial conditions affects only the probability of being clustered in one of the two states at $t = 2$, but not the transition probability matrix. A more parsimonious model could be then considered (see e.g. Bartolucci and Farcomeni, 2009), by allowing the initial conditions to change the location and the initial probabilities of the random effects only.

References.

- Adam, T., Griffiths, C., Leos-Barajas, V., Meese, E., Lowe, C., Blackwell, P., Righton, D. and Langrock, R. (2019). Joint modelling of multi-scale animal movement data using hierarchical hidden Markov models. *Methods in Ecology and Evolution*, 10, 1536–1550.
- Aitkin, M. and Alfó, M. (1998). Regression models for longitudinal binary responses. *Statistics and Computing*, 8, 289–307.
- Aitkin, M. (1999). A general maximum likelihood analysis of variance components in generalized linear models. *Biometrics*, 55, 117–128.
- Alfó, M. and Aitkin, M. (2000). Random coefficient models for binary longitudinal responses with attrition. *Statistics and Computing*, 10, 279–287.
- Aitkin, M. and Alfó, M. (2003). Longitudinal analysis of repeated binary data using autoregressive and random effects modelling. *Statistical Modelling*, 3, 291–303.
- Bartolucci, F., Farcomeni, A. and Pennoni, F. (2013). *Latent Markov model for longitudinal data*. CRC Press, Boca Raton.
- Bartolucci, F., Bacci, S. and Pennoni, F. (2014). Longitudinal analysis of self-reported health status by mixture latent autor-regressive models. *Journal of the Royal Statistical Society - Series C*, 63, 267–288.

TABLE 4

Estimates of the hidden Markov chain and mixture parameters for model M5.

Parameters		
$\boldsymbol{\pi} = \begin{bmatrix} \pi_1 \\ \pi_2 \end{bmatrix}$	$\begin{bmatrix} 0.692 \\ 0.308 \end{bmatrix}$	
	Initial conditions	
	$y_{n1} = 0$	$y_{n1} = 1$
$\boldsymbol{\delta} = \begin{bmatrix} \delta_1 \\ \delta_2 \end{bmatrix}$	$\begin{bmatrix} 0.833 \\ 0.167 \end{bmatrix}$	$\begin{bmatrix} 0.296 \\ 0.703 \end{bmatrix}$
$\mathbf{Q} = \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix}$	$\begin{bmatrix} 0.993 & 0.007 \\ 0.011 & 0.989 \end{bmatrix}$	$\begin{bmatrix} 1.000 & 0.000 \\ 0.020 & 0.980 \end{bmatrix}$

- Bartolucci, F. and Farcomeni, A. (2009). A multivariate extension of the dynamic logit model for longitudinal data based on a latent Markov heterogeneity structure. *Journal of the American Statistical Association*, 104, 816–831.
- Bartolucci, F. and Lupporelli, M. (2016). Pairwise likelihood inference for nested hidden Markov chain models for multilevel longitudinal data. *Journal of the American Statistical Association*, 111, 216–228.
- Baum, L.E., Petrie, T., Soules, G. and Weiss, N.A. (1970). Maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Annals of Mathematics and Statistics*, 41, 164–171.
- Bianconcini, S. and Bollen K.A. (2018) The Latent Variable-Autoregressive Latent Trajectory Model: A General Framework for Longitudinal Data Analysis. *Structural Equation Modeling: A Multidisciplinary Journal*, 25:5, 791–808.
- Bulla, J. and Berzel, A. (2008). Computational issues in parameter estimation for stationary hidden Markov models. *Computational Statistics*, 23, 1–18.
- DeRuiter, S.L., Langrock, R., Skributas, T., Goldbogen, J.A., Calambokidis, J., Friedlaender, A.S. and Southall, B.L. (2017). A multivariate mixed hidden Markov model for blue whale behaviour and responses to sound exposure. *The Annals of Applied Statistics*, 11, 362–392.
- Fotouhi, A. R. (2005). The initial conditions problem in longitudinal binary process: A simulation study. *Simulation Modelling Practice and Theory*, 13, 566–583.
- Früwirth-Schnatter, S., Pittner, S., Weber, A. and Winter-Ebmer, R. (2018). Analysing plant closure effects using time-varying mixture-of-experts Markov chain clustering. *Annals of Applied Statistics*, 12, 1796–1830.
- Gille, D. and Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. *The American Economic Review*, 102, 469–503.
- Hargreaves-Heap, S. and Varoufakis, Y. (2002). Some experimental evidence on the evolution of discrimination, co-operation and perceptions of fairness, *The Economic Journal*, 112, 679–703.
- Heckman, J. J. (1981). Heterogeneity and state dependence. In S. Rosen (ed.), *Studies in labor markets* (pp. 91139). Chicago, IL: Chicago University Press.

- Henry, J.F. (1999). John Locke, property rights and economic theory. *Journal of Economic Issues*, 33, 679–703.
- Lagona, F., Maruotti, A. and Padovano, F. (2015). Multilevel multivariate modelling of legislative count data, with a hidden Markov chain. *Journal of the Royal Statistical Society – Series A*, 178, 705–723.
- Leos-Barajas, V., Gangloff, E.J., Adam, T., Langrock, R., Morales, J.N., van Beest, F.M., and Nabe-Nielsen, J. (2017). Multi-scale modelling of animal movement and general behavior data using hidden Markov models with hierarchical structures. *Journal of Agricultural, Biological and Environmental Statistics*, 22, 232–248.
- Locke, J. (1980). *Second treatise of government*. Hackett Publishing.
- Maruotti, A. (2011). Mixed hidden Markov models for longitudinal data: an overview. *International Statistical Review*, 79, 427–454.
- Maruotti, A. and Rocci, R. (2012). A mixed non-homogeneous hidden Markov model for categorical data, with application to alcohol consumption. *Statistics in Medicine*, 31: 871–886.
- Maynard Smith, J. and Parker, G. A. (1976). The logic of asymmetric contests. *Animal behaviour*, 24(1), 159–175.
- Montanari, G.E., Doretto, M. and Bartolucci, F. (2018). A multilevel latent Markov model for the evaluation of nursing homes’ performance. *Biometrical Journal*, 60, 962–978.
- Pohle, J, Langrock, R., van Beest, F.M., and Schmidt, N.M. (2017). Selecting the number of states in hidden Markov models: pragmatic solutions illustrated using animal movement. *Journal of Agricultural, Biological and Environmental Statistics*, 22, 270–293.
- Rose, C. M. (2014). Psychologies of Property (and Why Property is Not a Hawk-Dove Game) in *The Philosophical Foundations of Property Law* Oxford University Press.
- Skrondal, A. and Rabe-Hesketh, S. (2014). Handling initial conditions and endogenous covariates in dynamic/transition models for binary data with unobserved heterogeneity. *Journal of the Royal Statistical Society, Series C*, 63, 211–237.
- Sugden, R. (1989). Spontaneous order. *The Journal of Economic Perspectives*, 3(4), 85–97.
- Sugden, R. (2004). *The economics of rights, co-operation and welfare*. Springer.
- Waldron, J. (2013). To Bestow Stability upon Possession: Humes Alternative to Locke. *Philosophical Foundations of Property Law*, 1–12.
- Welch, L.R. (2003). Hidden Markov models and the Baum-Welch algorithm. *IEEE Information Theory Society Newsletter*, 53, 1–13
- Wooldridge, J. M. (2005). Simple solutions to the initial conditions problem in dynamic, nonlinear panel data models with unobserved heterogeneity. *Journal of Applied Econometrics*, 20, 39–54.
- Zhang, S. and Chang, H.H. (to appear). A multilevel logistic hidden Markov model for learning under cognitive diagnosis. *Behavior Research Methods*, 52, 408–421.

ANTONELLO MARUOTTI
 DEPARTMENT OF MATHEMATICS
 UNIVERSITY OF BERGEN
 E-MAIL: antonello.maruotti@uib.it

ANTONELLO MARUOTTI
 DIPARTIMENTO DI GIURISPRUDENZA, ECONOMIA, POLITICA E
 DELLE LINGUE MODERNE
 LIBERA UNIVERSITÀ MARIA Ss. ASSUNTA
 E-MAIL: a.maruotti@lumsa.it

MARCO FABBRI
 GRADUATE SCHOOL OF ECONOMICS
 UNIVERSITY POMPEU FABRA
 E-MAIL: m.fabbri@uva.nl

MATTEO RIZZOLLI
 DIPARTIMENTO DI GIURISPRUDENZA, ECONOMIA, POLITICA E
 DELLE LINGUE MODERNE
 LIBERA UNIVERSITÀ MARIA Ss. ASSUNTA
 E-MAIL: m.rizzolli@lumsa.it