

Multivariate Behavioral Research

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/hmbr20

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

Antonello Maruotti, Marco Fabbri & Matteo Rizzolli

To cite this article: Antonello Maruotti, Marco Fabbri & Matteo Rizzolli (2021): Multilevel Hidden Markov Models for Behavioral Data: A Hawk-and-Dove Experiment, Multivariate Behavioral Research, DOI: 10.1080/00273171.2021.1912583

To link to this article: https://doi.org/10.1080/00273171.2021.1912583

Published online: 22 Jun 2021.

Submit your article to this journal

Article views: 8

View related articles

View Crossmark data

Full Terms & Conditions of access and use can be found at
https://www.tandfonline.com/action/journalInformation?journalCode=hmbr20

MULTIVARIATE BEHAVIORAL RESEARCH
https://doi.org/10.1080/00273171.2021.1912583

Routledge
Taylor & Francis Group

Check for updates

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

Antonello Maruotti[a,b], Marco Fabbri[c], and Matteo Rizzolli[b]

[a]Department of Mathematics, University of Bergen; [b]Dipartimento GEPLI, Libera Università Maria Ss. Assunta; [c]Department of Economics and Business, University Pompeu Fabra & Barcelona, Graduate School of Economics

ABSTRACT

Motivated by the analysis of behavioral data taken from an economic experiment based on the Hawk-and-Dove game, this article describes a multilevel hidden Markov model, that includes covariates, autoregression, and endogenous initial conditions under a unified framework. The data at hand are affected by multiple sources of latent heterogeneity, due to multilevel unobserved factors that operate in conjunction with observed covariates at all the levels of the data hierarchy. We fit a multilevel logistic regression model for repeated measurements of player behaviors, nested within groups of interacting players. The model integrates discrete random effects at the group level and Markovian sequences of discrete random effects at the player level. Parameters are estimated by a computationally feasible expectation-maximization algorithm. We model the probability of playing the Hawk strategy, which implies fighting aggressively for controlling an asset, and test the role played by initial possession, property, and other player-specific characteristics in driving hawkish behaviors. The results from our study suggest that crucial factors in determining hawkish behavior are both the way possession is achieved – which depends on our treatment manipulation- and possession itself. Furthermore, a clear time-dependence is observed in the data at the player level as accounted for by the Markovian random effects.

KEYWORDS

Hidden Markov models; initial conditions; conditional models; possession

1. Introduction

This article presents a model for the analysis of longitudinal hierarchical binary outcomes that extends standard random effects models to simultaneously account for serial dependence and heterogeneity at the different levels of the hierarchy. This particular model is used to study data collected from an economic laboratory experiment based on the Hawk-and-Dove game which is among the most famous games of evolutionary game theory as it can be used to explain the emergence of private property (Sugden, 2004). In the game, two players contend the control of an asset and they can either simultaneously choose between the *Hawk* strategy – i.e. fighting for controlling the asset, even if this might result in a negative payoff – or the *Dove* strategy – i.e. walking away from the fight if meeting an aggressive player, or sharing the asset (and incurring a reduction of the total resource value) if meeting another peaceful player. The particular payoff structure of this coordination game implies that the best strategy for a player confronting a peaceful Dove is to fight aggressively for the resource by playing Hawk. The strategy is also socially efficient, since the resource does not suffer from the value reduction following the Dove–Dove division. However, choosing the Hawk strategy when the game is played against another Hawk would result in the highest individual and social cost. In the game, costly fighting can be avoided and resources can be efficiently allocated when players coordinate so that at every encounter one consistently plays Hawk, while the other plays Dove or vice versa. How can this coordination be achieved? One way is by adopting the so called *bourgeois strategy*, which implies that the player who initially possesses the asset plays *Hawk* while the other player plays *Dove*. Although the *bourgeois strategy* is not necessarily the only theoretical solution resulting in an efficient equilibrium for the Hawk and Dove game, this coordination mechanism nicely explains how possession leads to the establishment of the property convention (Sugden, 1989), and how this convention can efficiently solve the coordination problem of costly fighting over the

CONTACT Antonello Maruotti a.maruotti@lumsa.it Department of Mathematics, University of Bergen, Bergen, Norway
© 2021 Taylor & Francis Group, LLC

appropriation of scarce resources (Maynard Smith & Parker, 1976; Sherratt & Mesterton-Gibbons, 2015).

In our lab experiment, groups of six players are divided in three couples and each one plays an Hawk-and-Dove game for the appropriation of an amount of tokens. The game is repeated for 10 periods and couples are rematched randomly at each period. We have multiple treatments in which we manipulate over two dimensions the way the tokens are initially presented to the players. The first treatment variation concerns whether the asymmetry between players introduced at the beginning of each round is based on possession or not. Specifically, in three out of five treatments one player in the couple is assigned possession of the tokens before playing the Hawk-and-Dove game. In the remaining two treatments, neither player possesses the tokens and players are differentiated by being associated to different colors. The second treatment variation concerns the procedure for the assignment of the initial possession/color, which can be determined by two different types of meritorious procedures (labor or first possession), or by simple luck. Combining these treatment manipulations, we observe whether and to which extent the convention of property – intended as the coordination of players on the bourgeois strategy – emerges.

By having both possessory and non-possessory color-based treatments, we are able to investigate whether possession is indeed a superior coordinating asymmetry. Moreover, meritoriousness may also play an important part in fostering the player's behavior. At last, the evolution of the hawkish behavior over time is investigated by looking at how the players learn to play from one round to another. With this design, we can address the following research questions:

- Does possession trigger the establishment of property? Does this favor an aggressive behavior?
- Does property induce more aggressive behaviors?
- Is merit important in establishing players' behavior?
- Is there a dependence over time in players behavior?

In the data the outcomes of players' decisions, i.e. to play hawk or dove, are measured as binary longitudinal hierarchical data: repeated measurements are collected for players 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 the research questions defined above.

As a general concern, when analyzing longitudinal hierarchical data, there are several specific features requiring attention in model building. In detail, the effects of one or more covariates on the response variable are often modeled in a regression framework, but also (observed and latent) time-dependence heterogeneity sources at the different levels of the hierarchy play a crucial role in data analysis, bearing in mind that the interest may also be on the association among repeated measurements. 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 probability of the occurrence of an event at the actual time is modeled as a function of lagged (past) values, i.e. the observed process has a first-order Markov property. The latter case, instead, implies the existence of heterogeneous subpopulations, whose behavior over time is generally different. 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 & Alfó, 1998, 2003; Alfó & Aitkin, 2000; Bianconcini & Bollen, 2018; Fotouhi, 2005; Skrondal & Rabe-Hesketh, 2014). To avoid misleading inference and interpretation, as a consequence of the inconsistency of the maximum likelihood estimators (Heckman, 1981), the dependence between the (endogenous) initial conditions and the random effects should be properly accounted for.

In detail, we develop a multilevel approach in which the unobserved heterogeneity at the different level of the hierarchy, i.e. unobserved players' and groups' behaviors, is assumed discrete and modeled through the inclusion of random effects (Aitkin, 1999; Bartolucci & Farcomeni, 2009). As for all discrete random effects models, the selection of the number of the support points is a part of the model selection process. This involves an additional burden, as models with a varying number of support points are fitted and compared in terms of the tradeoff between goodness of fit and parsimony. Nevertheless, the advantages of this approach are clear. Just to mention a few: the resulting model is a finite mixture (Aitkin, 1999) whose likelihood is computationally attractive and feasible; there is no need to specify any (unverifiable) parametric assumptions about the random effects distribution; the resulting components of the finite mixtures can be interpreted as typical behaviors. To account for the endogeneity of the initial conditions and conduct proper inference, we introduce a

Table 1. The Hawk-and-Dove game: payoffs of Player 1.

		Player 2	
		Hawk	Dove
Player 1	Hawk	−25	50
	Dove	0	15

modeling framework in a simple case of conditional likelihood of observations given the initial conditions. As a result, we allow for a change in the shape of the random effects distribution, 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 expectation-maximization (EM)-based algorithm (Baum et al., 1970). Different estimation methods and algorithms can be used to provide parameter estimates (Bulla & Berzel, 2008), also under the Bayesian paradigm (Rydén, 2008; Zhang & Berhane, 2014; Zhang & Chang, 2020).

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 follows. 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 modeling 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. Section 5 discusses the obtained results with respect to specific research questions and provides a comparison with other well-established approaches in a hidden Markov models setting. Conclusions and comments are given in Section 6.

2. The Hawk-and-Dove game and our experiment

The Hawk-and-Dove game was originally developed by theoretical biologists to model how animals interact when resources such as food or territory are scarce (Maynard Smith & Parker, 1976) and later, by analogy, it has also been used to study humans' coordination in the establishment and respect of property rights (Fabbri et al., 2021; Rose, 2013; Sugden, 1989, 2004; Waldron, 2013). Our study is situated within this particular stream of the literature.

Let us however recall the game in its original framing where two animals face a situation in which they can choose between playing the *Hawk* strategy or the *Dove* strategy. For the sake of simplicity, we here make direct reference to the payoffs used in our experiment. Hawks fight aggressively for an asset of value $V = 50$. However, if two Hawks meet, then the escalation of the fight produces a cost $C = -25$ for both Hawks. Doves retreats from the fight if faced with major escalation from an Hawk, thus obtaining a payoff of zero, or they share the asset at some cost $S = -10$ so that the overall payoff when a Dove meets another Dove is $50/2 - 10 = 15$, if it meets another Dove. Table 1 illustrates the payoffs matrix for Player 1. Player 2 payoffs' mirror the ones of Player 1. Notice that in evolutionary biology, the value of payoffs is measured in units of fitness while when playing with human participants, payoffs might be interpreted as measuring monetary units.

It is well known that the standard one-shot Hawk-and-Dove game has three Nash equilibria; two of them are in pure strategies (one player plays Hawk and the other player plays Dove and vice versa) and are efficient as aggregate payoffs are always maximal. There exists also a mixed strategy equilibrium where participants play Hawk with some positive probability which is instead inefficient because in some instances both players play Hawk or both players play Dove and this will result in nonmaximal payoffs. To understand the dynamic of this game, we turn to evolutionary game theory and assume the game is played repeatedly over time by players aiming at reproducing themselves as much as possible and where the best strategy is simply the one that over time survives natural selection. In this dynamic context, the payoff of each of the two strategies clearly depends on how many Hawks and Doves there are in the population. A well known result of this theoretical literature is that the only evolutionarily stable strategy (ESS) of the Hawk-and-Dove game is the inefficient Nash equilibrium in mixed strategies seen before that guarantees that both populations of Hawks and Dove coexist but also that costly fighting and inefficient sharing of resources happens repeatedly. This is the state of affairs unless some form of *asymmetry* between the two players, which should be payoff-irrelevant is introduced (Mesterton-Gibbons & Sherratt 2014). This asymmetry can be any characteristic that clearly introduces a distinction between players and assigns a "role" to each of the two players. For instance, it

Table 2. Main characteristics of the considered treatments.

Treatment	Type of information	Process of acquisition
Gift	Possessory	Arbitrary
Labor	Possessory	Meritorious
Treasure Trove	Possessory	Meritorious
Lucky Red	Colored	Arbitrary
Master Red	Colored	Meritorius

could be the distinction between whoever has the longer thumb (Eshel, 2005) or whoever is assigned a given colored label. *Possession* is such an asymmetry as one player may start the Hawk-and-Dove game initially controlling the asset, therefore playing with the role of *possessor* versus the other player being the *intruder*. Once such an asymmetry is present, the participants can exploit this additional information to coordinate on one of the two efficient equilibria; for instance, by establishing that the possessor plays Hawk, while the intruder plays Dove. Indeed, the strategy whereby the possessor plays Hawk while the intruder plays Dove is known in the literature as the *bourgeois* strategy and the standard theory shows that this strategy becomes an ESS and trumps the previous mixed strategy equilibrium (Sherratt & Mesterton-Gibbons, 2015). The introduction of the bourgeois strategy in the Hawk-and-Dove game provides a powerful and elegant model to stylize persistent observations of animal behavior such as the display of territoriality (Kokko et al., 2006) and the respect for other's property often observed in humans (Fabbri & Dari-Mattiacci, in press; Faillo et al., 2019). The Hawk-and-Dove game in fact shows that property behavior may emerge from the bottom up, as participants may simply find it efficient to respect each others' possessions rather than being the result of a top-down state driven imposition of third-party enforced property rights.

The Hawk-and-Dove game has been widely tested with experiments involving animals (Kokko et al., 2006), however there are surprisingly very few experiments involving human beings (Berninghaus et al., 2012; McAdams & Nadler, 2005; Neugebauer et al., 2008; Oprea et al., 2011). None of these articles studies the role of possession in explaining the choice of the Hawk strategy by players.

2.1. Experimental design

In our experiment, players interact in groups of six over ten rounds. In each round, three pairs are randomly matched and the matching is unknown to the players. Each player participates to two distinct activities. In the first part of each round, the asymmetry between players is introduced, whereas in the second part the two players play the Hawk-and-Dove game. The payoff matrix of the game is described in Table 1. The payoffs are expressed in Experimental Currency Units (ECU) and exchanged to euro at the end of the experiment at the rate of 50 ECU = 1 Euro. Participants begin the experiment with an initial endowment of roughly 300 ECU (the exact amount varies depending on the specific procedures of each treatment). As it is common in experimental economics, payoffs are calibrated so that the average payment to participants is roughly commensurable with the hourly wage of a nonspecialized worker. Indeed, experimental participants earned on average 14 Euro for an average duration of 70 min. One may wonder whether the results of the experiment may depend on the scale of payoffs. While it has been proven to be largely irrelevant for the experimental results in other paradigmatic games (Abeler et al., 2019; Larney et al., 2019) whether our results are robust to several-fold increases in the payoffs is an open question that only further experiments can definitely answer.

As said before, in each round there are two parts: in the first part the asymmetry between players is introduced whereas in the second part players play the Hawk-and-Dove game. Although this second part is the same for all sessions of the experiment, the first part varies and it is here where we introduce our treatment manipulation.

2.2. Experimental treatments and procedures

Our experimental manipulation occurs along two main dimensions. In the first dimension, we change whether the asymmetry is either *possessory* or *colored*. If there is possession, before starting one of the two players is given a claim to the asset and thus the other player is a potential intruder. In the non-possessory treatments instead, a different color is assigned to each of the two players. The second dimension concerns whether the assignment of roles (being the possessor or having a certain color) occurs as a result of a procedure based on luck (a random draw) or on merit (a pregame where players can display their skills). We thus define the asymmetry either *arbitrary* or *meritorious*.

In total there are five treatments, see Table 2. The possessory treatments mimic some ways of acquiring property typically found in property law such as *Gift, Labor*, and *Treasure Trove*.

In the *Gift* treatment, 50 tokens are randomly assigned to one player as windfall money while the

other player gets nothing. The asymmetry is thus possessory and arbitrary.

In the *Labor* treatment, participants undertake the same effort task as in the study by Gill and Prowse (2012): each participant must set the cursors of as many sliders as possible in a given time. In each pair, the player that correctly sets the largest number of cursors becomes the possessor of the 50 tokens. In this case, thus, the asymmetry is therefore possessory and meritorious. This activity, which does not require any particular ability or knowledge, is well suited for inducing participants to associate their possession to the intensity of their effort. In this way, this treatment mimics *labor* as a universally accepted mechanism of property acquisition (Henry, 1999; Locke, 1980).

In the *Treasure Trove* treatment, players take part to a treasure trove contest in order to gain control of 50 tokens before playing the Hawk-and-Dove game. Each player sees on screen the same 25 tiles which can be turned. Behind one of them a 50-token treasure is hidden and whoever registers the trove first gains possession. In this case, the asymmetry is thus again possessory and meritorious.

In the colored treatments, participants are no longer differentiated by being possessor/intruder; instead they are identified by either being *Red* or *Blue*, as in the study by Hargreaves-Heap and Varoufakis (2002),

In the *Lucky Red* treatment, the labels *Red* and *Blue* are assigned randomly at the beginning of each round. Notice that before the beginning of the Hawk-and-Dove experiment, nobody possesses the 50 tokens.

In the *Master Red* treatment, players take part to a treasure trove as in the Treasure Trove treatment. The difference lays in the fact that whoever registers the trove first gets the color Red (Blue) and not possession. Notice that before the beginning of the contest nobody possesses the 50 tokens.

We conducted different sessions in the lab involving 474 participants clustered into 79 groups. The experiment was digitalized using Z-tree (Fischbacher, 2007). Most of the participants were graduate and undergraduate students from LUISS University recruited via ORSEE (Greiner, 2015). Instructions were publicly announced to ensure common knowledge and control questions were casted to line up players' understanding of the experiment. A questionnaire on socio-demographic characteristics was collected at the end of the experiment. Participants were forbidden to communicate throughout the experiment. No missing values are reported.

2.3. Data description

A preliminary assessment of the behavior of the groups in playing hawk is described in Figure 1. We categorize the number of hawks in each group at each round into three categories to simplify data visualization: [0–1], [2–4], and [5–6] hawk-players. It is rather evident that the [0–1] category is rather rare, i.e. not many players decide to behave like doves. As expected, the [2–4] category is observed for large part of the groups and periods, i.e players are equally divided by those who play hawk and those who play dove. It is possible to identify groups whose behavior is rather constant over time; heterogeneity in groups' behaviors can also be inferred. 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 $[0-1], [2-4], [5-6]$. Persistence is not the norm for two of the three categories. If a small number of players ([0–1]) played hawk at time t, at time $t+1$ such a number would likely increase to [2–4] with *probability* 0.82 and, rarely, to [5–6] (only in the 2% of the cases). Similarly, if a large number of players played hawk at time t, it is likely that this number decreases at time $t+1$ passing to [2–4] or, rarely, to [0–1]. The mid category is the persistent one with the 76% of cases that the same number of aggressive behaviors is observed at two consecutive times. What is also rather clear is that it is more likely to pass from [2–4] to [5–6] rather than [0–1], as indication that the aggressive hawk behaviors are often preferred.

Aggressive behaviors may depend on several factors. A small increase in the overall aggressive behavior is observed, but significant differences are observed between possessor and intruder players (see Figure 2). Participants 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 intruder players the percentage of players playing hawk is also rather high, and above 50%, suggesting that, on average, the hawk behavior is preferred to the dove one.

We further expect that the asymmetry introduced by the treatments would also play a role in players' 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.

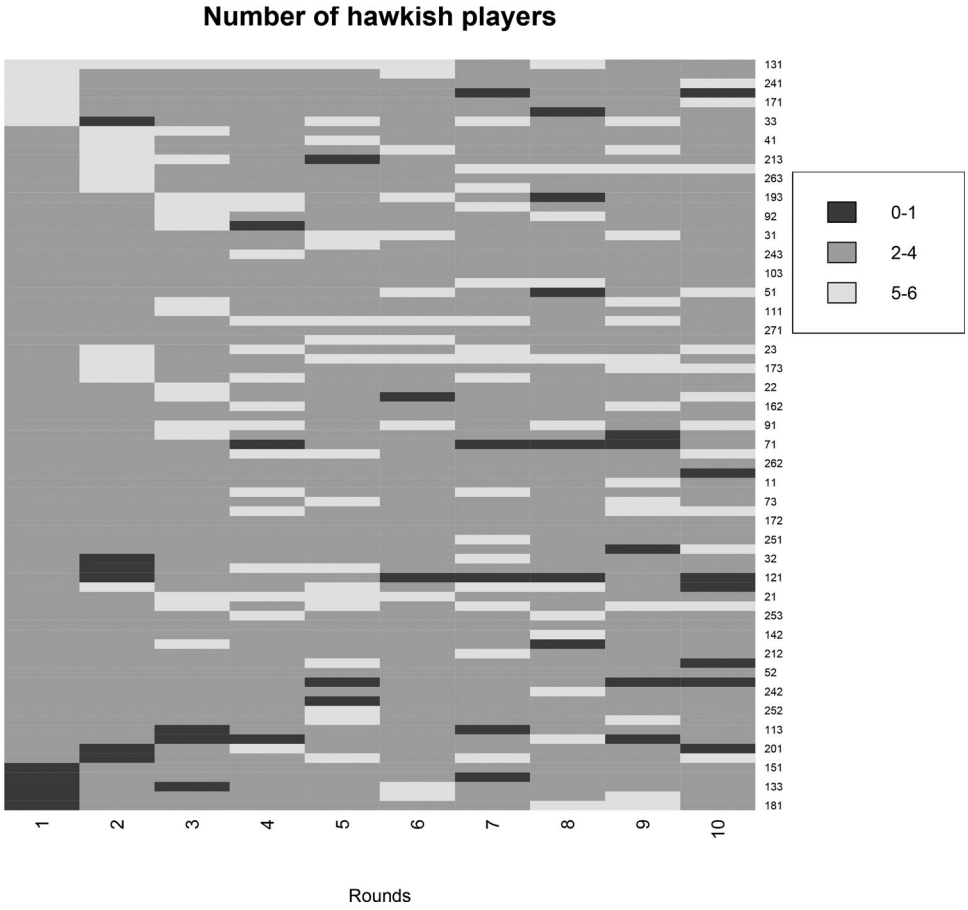


Figure 1. Number of players playing hawk over rounds per group.

3. Modeling

3.1. The multilevel hidden Markov model

The structure of the data considered in this article is as follows: we observe a certain number of players performing in an Hawk-and-Dove game, interacting within playing groups. A clear hierarchy can be easily detected in the data structure and different sources of (time-varying and time-constant) heterogeneity should be tackled at the different levels of the hierarchy.

Formally, the response variable, i.e. playing hawk, Y_{ntg} assumes two values only, {0,1}, and is observed on N $(n = 1, 2, ..., N_g)$ individuals clustered into G $(g = 1, 2, ..., G)$ groups at T_n $(t = 1, 2, ..., T_n)$ occasions. Since repeated measurements belonging to the same player assigned to a specific group are likely to be correlated, a modeling 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 player-specific covariates $\mathbf{x}_{ntg} = \{x_{ntg1}, ..., x_{ntgP}\}$, the responses are generated by combining individual- and group-specific random components. The individual-specific random component is specified by an array $\mathbf{b} = \{\mathbf{b}_1, ..., \mathbf{b}_N\}$ of N time-dependent trajectories, with $\mathbf{b}_n = \{b_{n1}, ..., b_{nT_n}\}$, 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 $p(\mathbf{v}) = p(v_1, ..., v_G)$.

A crucial assumption, common to most of the multilevel random effects models, is that the outcomes are conditionally independent given the players- and the group-specific random effects. Accordingly, the conditional distribution of $p(\mathbf{y}|\mathbf{b}, \mathbf{v})$ is given by

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

Under these assumptions, we have that

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

is fully specified by defining

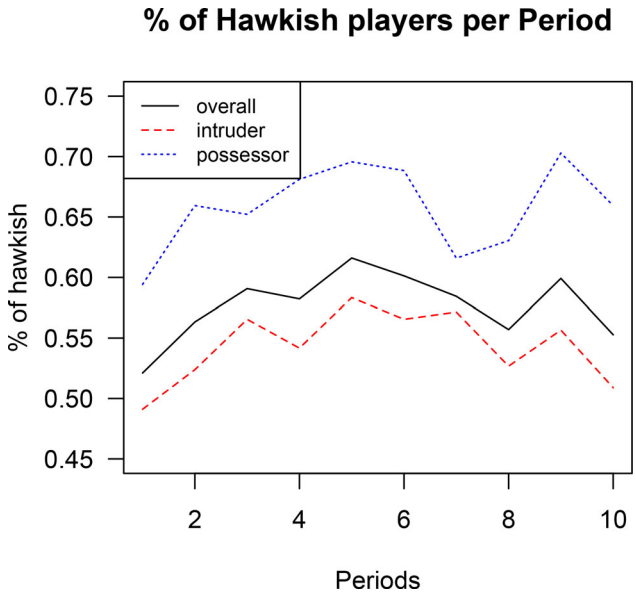


Figure 2. Percentage of players playing hawk over rounds, possessor vs. intruder players.

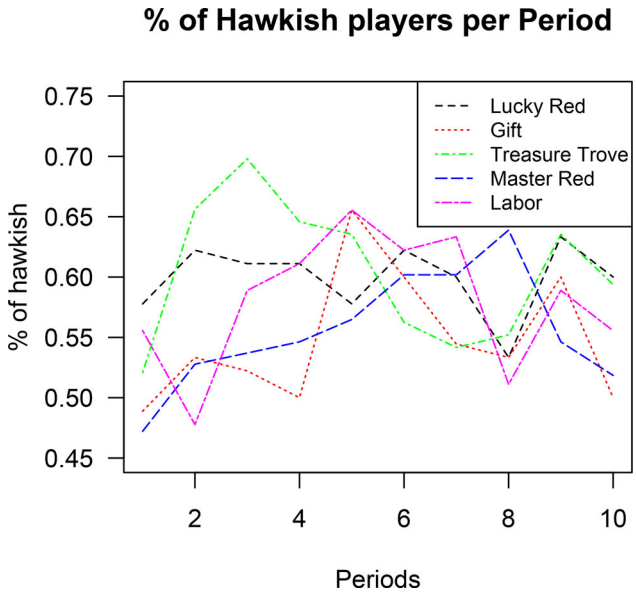


Figure 3. Percentage of players playing hawk over rounds by treatment.

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

$$\log\left(\frac{\lambda_{ntg}}{1-\lambda_{ntg}}\right) = \mathbf{x}'_{ntg}\boldsymbol{\beta} + \alpha y_{n(t-1)g} + v_g + b_{nt} \qquad (2)$$

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 allow for enough flexibility to model the distributions of the random terms with multiple modes and/or deviations for the Gaussian distribution (often taken for granted under the multilevel framework), we introduce discrete random effects, whose distributions have a finite number of mass points, in line with the nonparametric maximum likelihood principle (see e.g. Aitkin, 1999). Accordingly, the group-level random effects v_g are independently drawn for each group from a discrete distribution, say

$$v_g \sim p_M(\mathbf{v}; \boldsymbol{\pi}) \qquad (3)$$

where $p_M(\mathbf{v}; \boldsymbol{\pi})$ depends on a vector of M support points $\mathbf{v} = (v_1, ..., v_M)$ with mass probabilities $\boldsymbol{\pi} = (\pi_1, ..., \pi_M)$. At the player level, the random effects are time-varying and follow a first-order finite-state Markov chain with state-space $B = (b_1, ..., b_K)$. The parameters identifying the Markov chain are the initial distribution $\boldsymbol{\delta} = (\delta_1, ..., \delta_K), \delta_k = \Pr(b_{n1} = b_k)$, and a $K \times K$ transition probability matrix $\mathbf{Q}$ whose elements are given by $q_{nthk} = \Pr(b_{nt} = b_k | b_{n(t-1)} = b_h), t > 1; h, k = 1, 2, ..., K; \sum_{h=1}^{K} q_{nthk} = 1$. Homogeneous chains with $q_{nthk} = q_{hk}$ are the starting point of any analyses, but extensions to nonhomogeneous chains are trivial, and require minor efforts under our setting. Formally,

$$\mathbf{b}_n \sim p_K(\mathbf{b}; \boldsymbol{\delta}, \mathbf{Q}). \qquad (4)$$

Bearing in mind that this model can be also cast in the literature on the mixed hidden Markov models (Altman, 2007; de Haan-Rietdijk et al., 2017; Maruotti, 2011; Maruotti & Rydén, 2009) and is related to the approaches introduced by Marino and Alfó (2020); Bartolucci et al. (2014); Bartolucci and Lupparelli (2016); Leos-Barajas et al. (2017); DeRuiter et al. (2017); Montanari et al. (2018); Adam et al. (2019); Zhang and Chang (2020), this model includes some popular approaches to longitudinal data analysis as particular cases, see e.g. Lagona et al. (2015) for a brief discussion.

3.2. The initial conditions problem and the conditional approach

The joint distribution of the observed responses for player n, given the group-specific random effect v_g is given by

$$p(y_{n1g}, \ldots, y_{nT_ng}|\nu_g) = \sum_{\mathbf{b}_n} p_K(\mathbf{b}_n; \boldsymbol{\delta}, \mathbf{Q}) p(y_{n1g}|b_{n1}, \nu_g) \prod_{t=2}^{T_n} p(y_{ntg}|y_{n(t-1)g}, b_{nt}, \nu_g) \tag{5}$$

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

Nevertheless, Equation (5) is not determined from the model assumptions since the distribution $p(y_{n1g}|b_{n1}, \nu_g)$ is not specified from the model and the full likelihood is not available. Inference can be highly sensitive to misspecifcation of $p(y_{n1g}|b_{n1}, \nu_g)$. An alternative is to estimate regression parameters by maximizing the likelihood conditional on the first outcome Y_{n1g}, i.e. considering the distribution of $p(y_{n2g}, \ldots, y_{nTg}|y_{n1g})$:

$$p(y_{n2g}, \ldots, y_{nT_ng}|y_{n1g}, \nu_g) = \sum_{\mathbf{b}} p_K(\mathbf{b}_n, \boldsymbol{\delta}, \mathbf{Q}) \prod_{t=2}^{T_n} p(y_{ntg}|y_{n(t-1)g}, b_{nt}, \nu_g). \tag{6}$$

Bearing in mind that $\mathbf{b}_n$ is a player-specific effect shared by all player 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 modeled as a change in the location and shape of $\mathbf{b}_n$. Following Wooldridge (2005) and Aitkin and Alfó (2003), let us assume that

$$\mathbf{b}_n = \mathbf{b}_n^* + E[\mathbf{b}_n|y_{n1g}] \tag{7}$$

where

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

Then, Equation (2) can be rewritten as

$$\log\left(\frac{\lambda_{ntg}}{1-\lambda_{ntg}}\right) = \mathbf{x}'_{ntg}\boldsymbol{\beta} + \alpha y_{n(t-1)g} + \tilde{\alpha} y_{n1g} + \nu_g + b_{nt}. \tag{8}$$

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 player-specific random effects $\mathbf{b}_n$, i.e. we assume that Y_{n1g} affects the Markov chain parameters as well. Accordingly, we rewrite expression (6):

$$p(y_{n2g}, \ldots, y_{nT_ng}|y_{n1g}, \nu_g) = \sum_{\mathbf{b}} p_K(\mathbf{b}_n|y_{n1g}, \boldsymbol{\delta}_{y_{n1g}}, \mathbf{Q}_{y_{n1g}}) \prod_{t=2}^{T} p(y_{ntg}|y_{n(t-1)g}, y_{n1g}, b_{nt}, \nu_g). \tag{9}$$

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

$$\delta_{nk} = \Pr(b_{n2} = b_k|y_{n1g}) = \frac{\exp(\gamma_{0k} + \gamma_{1k} y_{n1g})}{\sum_{h=1}^{K} \exp(\gamma_{0h} + \gamma_{1h} y_{n1g})} \tag{10}$$

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

$$q_{nhk} = \Pr(b_{nt} = b_k|b_{n(t-1)} = b_h, y_{n1g}) = \frac{\exp(\phi_{0hk} + \phi_{1hk} y_{n1g})}{\sum_{k=1}^{K} \exp(\phi_{0hk} + \phi_{1hk} y_{n1g})} \tag{11}$$

Participants 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 Sections 3.1 and 3.2 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 locations of the time-constant group-specific random effects $(\nu_1, \ldots, \nu_m)$ and the related probabilities $\boldsymbol{\pi}_M$ of the random effects distribution at the group level;
c. the locations of the time-varying player-specific random effects $(b_1, \ldots, b_K)$ of the time-varying random effects;
d. the fixed effects $\boldsymbol{\gamma}$ and ϕ 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

$$L(\boldsymbol{\theta}|\mathbf{y}_1) = \prod_{g=1}^{G} \sum_{\mathbf{v}} p_M(\mathbf{v}; \boldsymbol{\pi}) \prod_{n=1}^{N_g} \sum_{\mathbf{b}} p_K(\mathbf{b}; \boldsymbol{\delta}_{y_{n1g}}, \mathbf{Q}_{y_{n1g}}) \prod_{t=2}^{T_n} p(y_{ntg}|y_{n(t-1)g}, y_{n1g}, b_{nt}, \nu_g; \boldsymbol{\beta}, \boldsymbol{\alpha}) \tag{12}$$

The conditional likelihood (12) can be efficiently computed by an extension of the forward recursion (Bartolucci et al., 2013; Zucchini et al., 2016). To find maximum likelihood estimates, we implement a version the EM algorithm, which has a long tradition of use in the hidden Markov models (HMM) literature; though numerical maximization could be used instead (MacDonald, 2014).

The algorithm is applied to the so-called complete-data log-likelihood function, where the player-specific hidden states and group-specific latent groups memberships are unknown, defined as

$$\ell_c(\boldsymbol{\theta}|\mathbf{y}_1) = \tag{13}$$
$$= \sum_{g=1}^{G}\sum_{m=1}^{M}\eta_{gm}\log(\pi_m)$$
$$+\sum_{n=1}^{N}\sum_{k=1}^{K}\xi_{n2k}\log(\delta_{nk}) \tag{14}$$
$$+\sum_{n=1}^{N}\sum_{k=1}^{K}\sum_{h=1}^{K}\sum_{t=3}^{T_n}\zeta_{nthk}\log(q_{nhk}) \tag{15}$$
$$+\sum_{g=1}^{G}\sum_{m=1}^{M}\sum_{k=1}^{K}\sum_{n=1}^{N}\sum_{t=2}^{T_n}\eta_{gm}\xi_{ntk}\log(p(y_{ntg}|y_{n(t-1)g}, y_{n1g}b_{nt}, \nu_g; \boldsymbol{\beta}, \boldsymbol{\alpha}) \tag{16}$$

where the variable $\eta_{gm} = I(\nu_g = \nu_m)$ 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 player 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 player n at time t belongs to state k and 0 otherwise.

At the E-step, the indicator variables η_{gm}, ξ_{ntk}, and ζ_{nthk} are replaced by their conditional expectations, given the observed data and the current value of the parameters. These conditional expectations can be straightforwardly obtained as a by-product of the forward–backward recursion following the procedure introduced in the mixed hidden Markov model literature (Maruotti, 2011). In the M-step, the conditional expected values of terms (13)–(16) are maximized separately. In particular, at iteration $r+1$, the maximum with respect to π_m has a closed form solution

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

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

$$\sum_{n=1}^{N}\sum_{k=1}^{K}\mathbb{E}(\xi_{n2k}|\mathbf{y}, \boldsymbol{\theta}^{(r)})\frac{\partial\log(\delta_{nk})}{\partial\boldsymbol{\gamma}_k} = 0 \tag{17}$$

and

$$\sum_{n=1}^{N}\sum_{t=3}^{T_n}\sum_{k=1}^{K}\mathbb{E}(\zeta_{nthk}|\mathbf{y}, \boldsymbol{\theta}^{(r)})\frac{\partial\log(q_{nhk})}{\partial\phi_h} = 0 \tag{18}$$

which are weighted sums of K equations with weights $\mathbb{E}(\xi_{n2k}|\mathbf{y}, \boldsymbol{\theta}^{(r)})$ and $\mathbb{E}(\zeta_{nthk}|\mathbf{y}, \boldsymbol{\theta}^{(r)})$, 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

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

The E- and M-steps are alternated repeatedly until the difference of the log-likelihoods at two consecutive iterations changes by an arbitrarily small amount that, in the current application, is 10^{-4}, as also suggested by Bulla and Berzel (2008).

4.1. Initialization of the algorithm, computational strategies, and model selection

The proposed multilevel HMMs share some common problems with several model-based clustering models. The likelihood equation will likely have multiple roots corresponding to local maxima, and the algorithms should be applied from a wide choice of starting values in any search for the maximum, aiming at excluding the so-called spurious maxima. Maruotti and Punzo (in press) discuss and critically evaluate several initialization strategies. Here, we follow a random start short-run strategy. In practice, the EM algorithm is run from a number of random initializations, and stopped without waiting for full convergence, i.e. when the log-likelihoods at two consecutive iterations differ less than 10^{-2}.Then the parameters corresponding to the maximum among these *short-run* likelihoods are used to initialize the *long-run* EM algorithm.

A major drawback of the EM algorithm concerns the computation of the uncertainty surrounding the estimates. Thus, to obtain standard errors, we consider a parametric bootstrap approach. The approximate standard error of each model parameter is then computed. This method can be summarized as follows. We fit the model, then generate a sample of observations from the fitted model (with the same sample size of the original data set and fixed covariates). The random effects are drawn from a discrete time-varying distribution at the player-level and

another discrete distribution at the group-level, according to the estimated random intercepts and the hidden/latent parameters. We then estimate the parameters for the bootstrap sample and repeat this procedure a large number of times (200 times in the considered empirical application) and record the values of the parameter estimates for each bootstrap sample. Afterward, we estimate the covariance matrix by the sample covariance matrix of the bootstrap estimates. Visser et al. (2000) investigate the reliable estimation of confidence bands in the context of HMMs and recommend bootstrap-based techniques, see also Clémençon et al. (2009).

Finally, we provide an idea of the computational burden required by our EM algorithm to achieve convergence. The code is written in R (R Core Team 2019). The elapsed time to fit a multilevel HMM with the lagged outcome, treatments, gender, age, total profit, and initial conditions as covariates, as described in Section 3.2, and player- and group-specific random effects is 127.76 s. Computation is performed on a Windows 10 PC, with Intel i7-1065G7 CPU 1.50 GHz, 16.0 GB RAM, using R 64 bit, and the elapsed time is computed via the proc.time() function of the base package.

A crucial aspect of the real data analysis deals with model selection, i.e. with the choice of the number of hidden states K and mixture components M. We fit models with a varying number of $K = \{2, 3, 4\}$ 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 Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) criteria (see e.g. Pohle et al., 2017, for a discussion).

To avoid the multinomial regressions defined by Equations (17) and (18), we could exploit the binary nature of the initial conditions, following the lines described by Maruotti (2015). Let $N_0 = \{n : Y_{n1g} = 0; n = 1, \ldots, N\}$ and $N_1 = \{n : Y_{n1g} = 1; n = 1, \ldots, N\}$ 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

$$L(\boldsymbol{\theta}|\mathbf{y}_1) = \prod_{g=1}^{G} \sum_{\mathbf{v}} p_M(\mathbf{v}; \boldsymbol{\pi})$$

$$\left\{ \prod_{n \in N_0} \sum_{\mathbf{b}} p_K(\mathbf{b}; \boldsymbol{\delta}_{y_{n1g}=0}, \mathbf{Q}_{y_{n1g}=0}) \prod_{t=2}^{T_n} p(y_{ntg}|y_{n(t-1)g}, y_{n1g}, b_{nt}, v_g; \boldsymbol{\beta}, \boldsymbol{\alpha}) \right.$$

$$\left. \prod_{n \in N_1} \sum_{\mathbf{b}} p_K(\mathbf{b}; \boldsymbol{\delta}_{y_{n1g}=1}, \mathbf{Q}_{y_{n1g}=1}) \prod_{t=2}^{T_n} p(y_{ntg}|y_{n(t-1)g}, y_{n1g}, b_{nt}, v_g; \boldsymbol{\beta}, \boldsymbol{\alpha}) \right\}.$$

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

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

and

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

5. Results

In this section, we illustrate the results obtained from the proposed approach to the dataset about players behavior in the Hawk-and-Dove game described in Section 2. We fitted, and compared, several models in terms of penalized likelihood criteria, namely:

- Model 1: a simple HMM, with no covariates;
- Model 2: an HMM with treatments, gender, age, and total profit as covariates;
- Model 3: an HMM with the lagged outcome, treatments, gender, age and total profit as covariates;
- Model 4: an HMM with the lagged outcome, treatments, gender, age, and total profit as covariates, and initial conditions included in the model as described in Section 3.2;
- Model 5: a multilevel HMM with the lagged outcome, treatments, gender, age, and total profit 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 3 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.

The estimates of the parameters under the selected models are reported in Table 4.

Regression parameters associated to the considered covariates are rather consistent across the models and

allow to draw some conclusions. First, 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 clearly affects behaviors. In particular, more aggressive behaviors arise when none of the players possess the tokens, i.e. under non-possessory treatments Lucky Red and Master Red rather than under possessory treatments. At the same time, owning the tokens is one of the main driving variables for playing hawk, i.e. also 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 possessor 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 & Parker, 1976): respect others' belongings and expect others to respect theirs. At the same time it is 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 *manna 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. Interestingly, the total amount of ECUs also plays a role. The highest the total profit the more likely the hawk strategy is. Thus, people are more likely to accept a defection and lose some ECUs, once they have some.

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

Table 3. Results from fitting different hidden Markov models with different values of *K*.

Model 1					
ℓ	*K*	*M*	No. parameters	AIC	BIC
−2,753.576	2		5	5,517.153	5,537.959
−2,706.514	3		10	5,433.641	5,474.641
−2,706.388	4		17	5,447.776	5,517.517
Model 2					
ℓ	*K*	*M*	No. parameters	AIC	BIC
−2,711.410	2		14	5,448.820	5,502.916
−2,659.211	3		18	5,354.421	5,429.323
−2,659.210	4		25	5,368.440	5,472.470
Model 3					
ℓ	*K*	*M*	No. parameters	AIC	BIC
−2,706.880	2		14	5,441.761	5,500.018
−2,659.600	3		19	5,357.200	5,436.263
−2,659.430	4		26	5,370.860	5,479.051
Model 4					
ℓ	*K*	*M*	No. parameters	AIC	BIC
−2,611.315	2		17	5,256.631	5,327.371
−2,603.098	3		28	5,262.196	5,378.710
−2,599.846	4		43	5,285.693	5,464.625
Model 5					
ℓ	*K*	*M*	No. parameters	AIC	BIC
−2,281.686	2	2	20	**4,603.372**	**4,686.596**
−2,277.936	3	2	32	4,619.872	4,753.031
−2,277.935	4	2	49	4,651.870	4,851.608

The maximum log-likelihood of each model is denoted by ℓ. Bold values indicates the lowest AIC and BIC values.

Table 4. Regression estimates from fitting different hidden Markov models (standard errors in brackets).

Variable/Parameter	Model 1	Model 2	Model 3	Model 4	Model 5
Fixed effects					
Intercept (β_0)	0.604 (0.164)	0.193 (0.330)	0.166 (0.335)	−1.928 (0.377)	−2.186 (0.397)
Labor	Reference treatment				
Lucky Red (β_1)	–	0.385 (0.169)	0.395 (0.171)	0.406 (0.121)	0.465 (0.122)
Gift (β_2)	–	−0.194 (0.184)	−0.202 (0.186)	−0.201 (0.114)	−0.175 (0.114)
Treasure Trove (β_3)	–	0.162 (0.172)	0.168 (0.175)	0.147 (0.116)	0.150 (0.117)
Master Red (β_4)	–	0.319 (0.168)	0.324 (0.170)	0.350 (0.118)	0.370 (0.118)
Possessor (β_5)	–	0.884(0.104)	0.893 (0.105)	0.833 (0.093)	0.825 (0.093)
Male (β_6)	–	0.016 (0.108)	0.014 (0.110)	−0.030 (0.071)	0.010 (0.071)
Age (β_7)	–	−0.002 (0.023)	−0.002 (0.024)	0.022 (0.016)	0.027 (0.016)
ECUs (β_8)	–	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)
Lag-outcome (α)	–	–	−0.046 (0.100)	0.069 (0.073)	0.070 (0.038)
Initial conditions ($\tilde{\alpha}$)	–	–	–	2.685 (0.136)	2.852 (0.146)
Subject-specific random effects					
$b_{nt} = b_1$	−2.390	−2.726	−2.801	−1.394	−1.438
$b_{nt} = b_2$	−0.272	−0.268	−0.281	1.245	1.293
$b_{nt} = b_3$	3.345	3.095	3.088	–	–
Group-specific random effects					
$v_g = v_1$	–	–	–	–	−0.165
$v_g = v_2$	–	–	–	–	0.273

Table 5. Estimates of the hidden Markov chain and mixture parameters for Model5.

Parameters			
$\boldsymbol{\pi} = \begin{bmatrix} \pi_1 \\ \pi_2 \end{bmatrix}$	$\begin{bmatrix} 0.624 \\ 0.376 \end{bmatrix}$		
	Initial conditions		
	$y_{n1} = 0$	$y_{n1} = 1$	
$\boldsymbol{\delta} = \begin{bmatrix} \delta_1 \\ \delta_2 \end{bmatrix}$	$\begin{bmatrix} 0.149 \\ 0.851 \end{bmatrix}$	$\begin{bmatrix} 0.683 \\ 0.317 \end{bmatrix}$	
$\mathbf{Q} = \begin{bmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{bmatrix}$	$\begin{bmatrix} 0.999 & 0.001 \\ 0.010 & 0.990 \end{bmatrix}$	$\begin{bmatrix} 1.000 & 0.000 \\ 0.034 & 0.966 \end{bmatrix}$	

subsequent times, indicating persistence of behaviors, not due to the way players 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 Model 5. For the multilevel model, Model 5, estimates of the hidden Markov chain and mixture parameters are reported in Table 5. Two states and two components are estimated at the player and group levels, respectively. The locations of the state- and component-specific random effects represent the propensity of playing hawk at the player and group levels, respectively. States represent two different players, more or less aggressive, behaviors, and mixture component different ways group coordination. In State 1 ($\mathbf{b}_{nt} = b_1 = -1.438$), players are less likely to play hawk, while State 2 ($\mathbf{b}_{nt} = b_2 = 1.293$) clusters more aggressive attitudes. The two states are well separated, and as expected they communicate rarely (see the transition matrices in Table 5). So the behaviors at the player level are well identified and characterized by strong persistence over time. The group effects have lower impacts on the probability of playing hawk, as the estimated locations $v_g = v_1 = -0.165$ and $v_g = v_2 = 0.273$ have smaller magnitude than their player counterpart. It deserves to be mentioned that around $\pi_1 \times 100 = 62.4\%$ of the groups (i.e. those clustered in the first mixture component) show a propensity to sharing, further reducing the probability of playing hawk. 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 affect 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 & Farcomeni, 2009), by allowing the initial conditions to change the location and the initial probabilities of the random effects only.

6. Conclusions

In this work, we introduce a multilevel hidden Markov model to handle different sources of heterogeneity in a Hawk-and-Dove longitudinal experiment. The proposed approach is able to distinguish between player- and group-specific heterogeneity accounting for the dynamics of the players behavior considering an autoregressive model nested into a hidden Markov chain, disentangling between true and spurious contagion. The parameters are estimated in a maximum likelihood framework and the implementation of the EM algorithm is not cumbersome, though it involves a large number of quantities to be estimated.

From a pure empirical perspective, the Hawk–Dove game has received very limited attention by experimental researchers and, to our knowledge, it has never been studied in the lab by introducing asymmetric information on possession. In our experiment, we manipulate on the one hand the type of asymmetry provided and we, further, manipulate the process of acquisition, that is to say whether the role of possessor player is acquired arbitrarily or meritoriously.

Our experimental results show that hawk behaviors are likely to be player-driven rather than group-driven and can be affected by treatment manipulation. Differences are estimated in the probability of playing hawk for the different treatments, but observed and unobserved player-specific characteristics play a major role.

Of course the proposed model could be extended allowing for time-varying heterogeneity at the group level, but the results show that the impact of groups on the probability of playing hawk is rather marginal, given the player characteristics. Moreover, including a further Markov chain at the group level would considerably increase the computational burden involved in the estimation step and requires more sophisticated and time-consuming algorithms (see e.g. Bartolucci & Lupparelli, 2016). Similarly, we may allow some or all regression parameters to vary across states, increasing the flexibility of the model at the price of a higher number of parameters to be estimated.

A crucial implicit assumption of the proposed approach relates to the sojourn distribution, i.e. the time spent in each hidden state; by construction, the sojourn distribution is geometric. This assumption could be easily relaxed allowing for an arbitrary sojourn distribution, parametrizing the model in the fashion of Langrock and Zucchini (2011). The proposed approach is rather flexible and accounts for several data features, but it somehow belongs to the family of generalized linear mixed models. This also

could be extended allowing for hidden Markov generalized additive models (Langrock et al., 2017).

We would like to further remark that the proposed approach can be easily extended to deal with missing values. If the missing process is at random, minor modifications of the likelihood computation are required, see e.g. Section 2.3.4 in Zucchini et al. (2016) for the time series case. Dealing with missing-not-at-random values is a bit more complex. Several approaches are proposed in the literature (Bartolucci & Farcomeni, 2019; Cai et al., 2018; Marino et al., 2018; Marino & Alfó, 2020). The pattern mixture, conditional approach by Maruotti (2015) can be easily cast in the current setting with minor efforts.

At last, as noticed by a reviewer, whereas the model proposed assumes that observations are conditionally independent given the group-level and subject-level random effects, observations are collected in pairs (three pairs of participants per group). Dyads constitute a special case of hierarchically structured data with variation at both the individual and dyadic level. Therefore, certain degree of correlation could be expected between the observations of dyads of participants playing the game at a given time t. This would represent a deviation from the assumption of independence of observation given the groups' and the subjects' random effects. McMahon et al. (2006); Kenny and Kashy (2010); Wickham and Macia (2019) discuss multilevel models including a dyad random effect too. Our analysis ignores dyad effects and do not account for inter-dyad correlation; as a possible consequence, the variability of the regression parameters may be understated. This should be investigated with an ad-hoc Monte Carlo simulation study.

Article information

Conflict of interest disclosures: Each author signed a form for disclosure of potential conflicts of interest. No authors reported any financial or other conflicts of interest in relation to the work described.

Ethical Principles: The authors affirm having followed professional ethical guidelines in preparing this work. These guidelines include obtaining informed consent from human participants, maintaining ethical treatment and respect for the rights of human or animal participants, and ensuring the privacy of participants and their data, such as ensuring that individual participants cannot be identified in reported results or from publicly available original or archival data.

Funding: The research of Marco Fabbri and Matteo Rizzolli has been partially supported by the Einaudi Institute for Economics and Finance Research Grant: "*Being Bourgeois: Cross-cultural Experiments on the Establishment of Private Property*". The research of Antonello Maruotti has been partially supported by the Financial Market Fund (Norwegian Research Council), project number 309218 "Statistical modelling and inference for (high-dimensional) financial data."

Role of the Funders/Sponsors: None of the funders or sponsors of this research had any role in the design and conduct of the study; collection, management, analysis, and interpretation of data; preparation, review, or approval of the manuscript; or decision to submit the manuscript for publication.

Acknowledgments: The authors would like to thank the Associate Editor and two anonymous reviewers for their valuable comments. The ideas and opinions expressed herein are those of the authors alone, and endorsement by the authors' institutions is not intended and should not be inferred.

References

Abeler, J., Nosenzo, D., & Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, *87*(4), 1115–1153. https://doi.org/10.3982/ECTA14673

Adam, T., Griffiths, C., Leos-Barajas, V., Meese, E., Lowe, C., Blackwell, P., Righton, D., & Langrock, R. (2019). Joint modelling of multi-scale animal movement data using hierarchical hidden Markov models. *Methods in Ecology and Evolution*, *10*(9), 1536–1550. https://doi.org/10.1111/2041-210X.13241

Aitkin, M. (1999). A general maximum likelihood analysis of variance components in generalized linear models. *Biometrics*, *55*(1), 117–128. https://doi.org/10.1111/j.0006-341x.1999.00117.x

Aitkin, M., & Alfó, M. (1998). Regression models for longitudinal binary responses. *Statistics and Computing*, *8*(4), 289–307. https://doi.org/10.1023/A:1008847820371

Aitkin, M., & Alfó, M. (2003). Longitudinal analysis of repeated binary data using autoregressive and random effects modelling. *Statistical Modelling*, *3*(4), 291–303. https://doi.org/10.1191/1471082X03st061oa

Alfó, M., & Aitkin, M. (2000). Random coefficient models for binary longitudinal responses with attrition. *Statistics and Computing*, *10*(4), 279–287. https://doi.org/10.1023/A:1008999824193

Altman, R. J. (2007). Mixed hidden Markov models: An extension of the hidden Markov model to the longitudinal data setting. *Journal of the American Statistical Association*, *102*(477), 201–210. https://doi.org/10.1198/016214506000001086

Bartolucci, F., Bacci, S., & Pennoni, F. (2014). Longitudinal analysis of self-reported health status by mixture latent autor-regressive models. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, *63*(2), 267–288. https://doi.org/10.1111/rssc.12030

Bartolucci, F., & 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*(486), 816–831. https://doi.org/10.1198/jasa.2009.0107

Bartolucci, F., & Farcomeni, A. (2019). A shared-parameter continuous-time hidden Markov and survival model for

longitudinal data with informative dropout. *Statistics in Medicine*, *38*(6), 1056–1073. https://doi.org/10.1002/sim.7994

Bartolucci, F., Farcomeni, A., & Pennoni, F. (2013). *Latent Markov model for longitudinal data*. CRC Press.

Bartolucci, F., & Lupparelli, M. (2016). Pairwise likelihood inference for nested hidden Markov chain models for multilevel longitudinal data. *Journal of the American Statistical Association*, *111*(513), 216–228. https://doi.org/10.1080/01621459.2014.998935

Baum, L. E., Petrie, T., Soules, G., & Weiss, N. A. (1970). Maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *The Annals of Mathematical Statistics*, *41*(1), 164–171. https://doi.org/10.1214/aoms/1177697196

Berninghaus, S. K., Ehrhart, K. M., & Ott, M. (2012). Forward-looking behavior in Hawk–Dove games in endogenous networks: Experimental evidence. *Games and Economic Behavior*, *75*(1), 35–52. https://doi.org/10.1016/j.geb.2011.12.003

Bianconcini, S., & Bollen, K. A. (2018). The latent variable-autoregressive latent trajectory model: A general framework for longitudinal data analysis. *Structural Equation Modeling*, *25*(5), 791–808. https://doi.org/10.1080/10705511.2018.1426467

Bulla, J., & Berzel, A. (2008). Computational issues in parameter estimation for stationary hidden Markov models. *Computational Statistics*, *23*(1), 1–18. https://doi.org/10.1007/s00180-007-0063-y

Cai, J., Ouyang, M., Kang, K., & Song, X. (2018). Bayesian diagnostics of hidden Markov structural equation models with missing data. *Multivariate Behavioral Research*, *53*(2), 151–171. https://doi.org/10.1080/00273171.2017.1407233

Clémençon, S., Garivier, A., & Tressou, J. (2009). *Pseudo-regenerative block-bootstrap for hidden markov chains* [Paper presentation]. 2009 IEEE/SP 15th Workshop on Statistical Signal Processing (pp. 465–468). Cardiff, UK.

de Haan-Rietdijk, S., Kuppens, P., Bergeman, C. S., Sheeber, L. B., Allen, N. B., & Hamaker, E. L. (2017). On the use of mixed Markov models for intensive longitudinal data. *Multivariate Behavioral Research*, *52*(6), 747–767. https://doi.org/10.1080/00273171.2017.1370364

DeRuiter, S. L., Langrock, R., Skirbutas, T., Goldbogen, J. A., Calambokidis, J., Friedlaender, A. S., & 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*(1), 362–392. https://doi.org/10.1214/16-AOAS1008

Eshel, I. (2005). Asymmetric population games and the legacy of Maynard Smith: From evolution to game theory and back? *Theoretical Population Biology*, *68*(1), 11–18. https://doi.org/10.1016/j.tpb.2004.11.003

Fabbri, M., & Dari-Mattiacci, G. (in press). The virtuous cycle of property. *Review of Economics and Statistics*. https://doi.org/10.1162/rest_a_00905.

Fabbri, M., Rizzolli, M., & Maruotti, A. (2021). Possession is Nine-Tenths of the Law: possession, property, and coordination in a Hawk–Dove Experiment. *Journal of Institutional Economics*, *17*(2), 267–288. https://doi.org/10.1017/S1744137420000442

Faillo, M., Rizzolli, M., & Tontrup, S. (2019). Thou shalt not steal: Taking aversion with legal property claims. *Journal of Economic Psychology*, *71*, 88–101. https://doi.org/10.1016/j.joep.2018.08.009

Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178. https://doi.org/10.1007/s10683-006-9159-4

Fotouhi, A. R. (2005). The initial conditions problem in longitudinal binary process: A simulation study. *Simulation Modelling Practice and Theory*, *13*(7), 566–583. https://doi.org/10.1016/j.simpat.2005.02.004

Gill, D., & Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review*, *102*(1), 469–503. https://doi.org/10.1257/aer.102.1.469

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, *1*(1), 114–125. https://doi.org/10.1007/s40881-015-0004-4

Hargreaves-Heap, S., & Varoufakis, Y. (2002). Some experimental evidence on the evolution of discrimination, cooperation and perctions of fairness. *The Economic Journal*, *112*(481), 679–703. https://doi.org/10.1111/1468-0297.00735

Heckman, J. J. (1981). Heterogeneity and state dependence. In S. Rosen (Ed.), *Studies in labor markets* (pp. 91–139). Chicago University Press.

Henry, J. F. (1999). John Locke, property rights and economic theory. *Journal of Economic Issues*, *33*(3), 609–703. https://doi.org/10.1080/00213624.1999.11506188

Kenny, D. A., & Kashy, D. A. (2010). Dyadic data analysis using multilevel modeling. In Joop Hox & J. Kyle Roberts (Eds.), *Handbook of advanced multilevel analysis* (pp. 335–370). Routledge.

Kokko, H., López-Sepulcre, A., & Morrell, L. J. (2006). From hawks and doves to self-consistent games of territorial behavior. *The American Naturalist*, *167*(6), 901–912. https://doi.org/10.1086/504604

Lagona, F., Maruotti, A., & Padovano, F. (2015). Multilevel multivariate modelling of legislative count data, with a hidden Markov chain. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, *178*(3), 705–723. https://doi.org/10.1111/rssa.12089

Langrock, R., Kneib, T., Glennie, R., & Michelot, T. (2017). Markov-switching generalized additive models. *Statistics and Computing*, *27*(1), 259–270. https://doi.org/10.1007/s11222-015-9620-3

Langrock, R., & Zucchini, W. (2011). Hidden Markov models with arbitrary state dwell-time distributions. *Computational Statistics & Data Analysis*, *55*(1), 715–724. https://doi.org/10.1016/j.csda.2010.06.015

Larney, A., Rotella, A., & Barclay, P. (2019). Stake size effects in ultimatum game and dictator game offers: A meta-analysis. *Organizational Behavior and Human Decision Processes*, *151*, 61–72. https://doi.org/10.1016/j.obhdp.2019.01.002

Leos-Barajas, V., Gangloff, E. J., Adam, T., Langrock, R., Morales, J. N., van Beest, F. M., & 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*(3), 232–248. https://doi.org/10.1007/s13253-017-0282-9

Locke, J. (1980). *Second treatise of government*. Hackett Publishing.

MacDonald, I. L. (2014). Numerical maximisation of likelihood: A neglected alternative to EM? *International Statistical Review*, *82*(2), 296–308. https://doi.org/10.1111/insr.12041

Marino, F., & Alfó, M. (2020). Finite mixtures of hidden Markov models for longitudinal responses subject to drop out. *Multivariate Behavioral Research*, *55*(5), 647–663. https://doi.org/10.1080/00273171.2019.1660606

Marino, M. F., Tzavidis, N., & Alfò, M. (2018). Mixed hidden Markov quantile regression models for longitudinal data with possibly incomplete sequences. *Statistical Methods in Medical Research*, *27*(7), 2231–2246. https://doi.org/10.1177/0962280216678433

Maruotti, A. (2011). Mixed hidden Markov models for longitudinal data: an overview. *International Statistical Review*, *79*(3), 427–454. https://doi.org/10.1111/j.1751-5823.2011.00160.x

Maruotti, A. (2015). Handling non-ignorable dropouts in longitudinal data: a conditional model based on a latent Markov heterogeneity structure. *TEST*, *24*(1), 84–109. https://doi.org/10.1007/s11749-014-0397-z

Maruotti, A., & Punzo, A. (in press). Initialization of hidden Markov and semi-Markov models: a critical evaluation of several strategies. *International Statistical Review*. https://doi.org/10.1111/insr.12436.

Maruotti, A., & Rocci, R. (2012). A mixed non-homogeneous hidden Markov model for categorical data, with application to alcohol consumption. *Statistics in Medicine*, *31*(9), 871–886. https://doi.org/10.1002/sim.4478

Maruotti, A., & Rydén, T. (2009). A semiparametric approach to hidden Markov models under longitudinal observations. *Statistics and Computing*, *19*(4), 381–393. https://doi.org/10.1007/s11222-008-9099-2

Maynard Smith, J., & Parker, G. A. (1976). The logic of asymmetric contests. *Animal Behaviour*, *24*(1), 159–175. https://doi.org/10.1016/S0003-3472(76)80110-8

McAdams, R. H., & Nadler, J. (2005). Testing the focal point theory of legal compliance: the effect of third-party expression in an experimental hawk/dove game. *Journal of Empirical Legal Studies*, *2*(1), 87–123. https://doi.org/10.1111/j.1740-1461.2005.00032.x

McMahon, J. M., Pouget, E. R., & Tortu, S. (2006). A guide for multilevel modeling of dyadic data with binary outcomes using SAS PROC NLMIXED. *Computational Statistics & Data Analysis*, *50*(12), 3663–3680. https://doi.org/10.1016/j.csda.2005.08.008

Mesterton-Gibbons, M., & Sherratt, T. N. (2014). Bourgeois versus anti-Bourgeois: A model of infinite regress. *Animal Behaviour*, *89*, 171–183. https://doi.org/10.1016/j.anbehav.2014.01.002

Montanari, G. E., Doretti, M., & Bartolucci, F. (2018). A multilevel latent Markov model for the evaluation of nursing homes' performance. *Biometrical Journal. Biometrische Zeitschrift*, *60*(5), 962–978. https://doi.org/10.1002/bimj.201700223

Neugebauer, T., Poulsen, A., & Schram, A. (2008). Fairness and reciprocity in the hawk-dove game. *Journal of Economic Behavior & Organization*, *66*(2), 243–250. https://doi.org/10.1016/j.jebo.2005.12.008

Oprea, R., Henwood, K., & Friedman, D. (2011). Separating the Hawks from the Doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, *146*(6), 2206–2225. https://doi.org/10.1016/j.jet.2011.10.014

Pohle, J., Langrock, R., van Beest, F. M., & 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*(3), 270–293. https://doi.org/10.1007/s13253-017-0283-8

R Core Team. (2019). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.

Rose, C. M. (2013). Psychologies of property (and why property is not a Hawk-Dove game). In James Penner & Henry Smith (Eds.), *The philosophical foundations of property law* (pp. 13–20). Oxford University Press.

Rydén, T. (2008). EM versus Markov chain Monte Carlo for estimation of hidden Markov models: A computational perspective. *Bayesian Analysis*, *3*(4), 659–688. https://doi.org/10.1214/08-BA326

Sherratt, T. N., & Mesterton-Gibbons, M. (2015). The evolution of respect for property. *Journal of Evolutionary Biology*, *28*(6), 1185–1202. https://doi.org/10.1111/jeb.12648

Skrondal, A., & 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 (Applied Statistics)*, *63*(2), 211–237. https://doi.org/10.1111/rssc.12023

Sugden, R. (1989). Spontaneous order. *Journal of Economic Perspectives*, *3*(4), 85–97. https://doi.org/10.1257/jep.3.4.85

Sugden, R. (2004). *The economics of rights, co-operation and welfare*. Palgrave Macmillan.

Visser, I., Raijmakers, M. E. J., & Molenaar, P. C. M. (2000). Confidence intervals for hidden Markov model parameters. *British Journal of Mathematical and Statistical Psychology*, *53*(2), 317–327. https://doi.org/10.1348/000711000159240

Waldron, J. (2013). To bestow stability upon possession: Hume's alternative to locke. In James Penner & Henry Smith (Eds.), *Philosophical foundations of property law* (pp. 1–12). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199673582.003.0001.

Wickham, R. E., & Macia, K. S. (2019). Examining cross-level effects in dyadic analysis: A structural equation modeling perspective. *Behavior Research Methods*, *51*(6), 2629–2645. https://doi.org/10.3758/s13428-018-1117-5

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*(1), 39–54. https://doi.org/10.1002/jae.770

Zhang, S., & Chang, H. H. (2020). A multilevel logistic hidden Markov model for learning under cognitive diagnosis. *Behavior Research Methods*, *52*(1), 408–421. https://doi.org/10.3758/s13428-019-01238-w

Zhang, Y., & Berhane, K. (2014). Bayesian mixed hidden Markov models: A multi-level approach to modeling categorical outcomes with differential misclassification. *Statistics in Medicine*, *33*(8), 1395–1408. https://doi.org/10.1002/sim.6039

Zucchini, W., MacDonald, I. L., & Langrock, R. (2016). *Hidden Markov models for time series: An introduction using R* (2nd ed.). Chapman & Hall.