Graph-Based Modeling of the Multi-Agent Scheduling Problem

Prunelle Vogler [*] Léa Ricard [*] Frédéric Meunier [†]
Michel Bierlaire [*]

December 20, 2025

Report TRANSP-OR 251220
Transport and Mobility Laboratory
School of Architecture, Civil and Environmental Engineering
École polytechnique fédérale de Lausanne
transp-or.epfl.ch

[*]École polytechnique fédérale de Lausanne (EPFL), School of Architecture, Civil and Environmental Engineering (ENAC), Transport and Mobility Laboratory, Switzerland, @epfl.ch

[†]École nationale des ponts et chaussées, CERMICS, @enpc.fr

# Abstract

This paper introduces a new graph-based mixed-integer linear formulation for the multi-agent scheduling problem arising in activity-based models. The proposed framework generalizes existing approaches, notably that of Rezvany et al. (2023), by modeling activities, travel, and coordination decisions within a unified labeled graph structure.

Each vertex represents an activity performed by a subset of agents at a specific location, while arcs capture travel and transport mode choices. This representation enforces synchronization compactly, eliminates redundant variables, and yields interpretable solutions, as each agent's schedule corresponds to a distinct path in the graph.

The formulation substantially improves computational performance compared to previous models and provides a flexible foundation for representing new behavioral situations. Shared activities are represented explicitly as vertices where the paths of multiple agents meet, making coordination constraints both intuitive and compact. Additional features—such as mandatory activity groups, participation limits, or shared resources—can be incorporated with minimal effort.

Numerical experiments on real household data confirm the computational gains, showing an average speed-up of almost one order of magnitude compared to the previous formulation, while capturing a richer variety of coordination behaviors. These results illustrate the potential of the proposed model for next-generation activity-based modeling.

Keywords: Activity-Based Modeling, Mixed-Integer Programming, Multi-Agent Scheduling.

1 Introduction

Activity-based models (ABMs) have gained significant attention in transportation and urban planning due to their ability to realistically capture how individuals plan their activities and travels over time. Unlike traditional trip-based models, ABMs account for the dynamic decision-making processes of individuals, taking into account constraints like time, resources, and interpersonal interactions. These models have been particularly useful for understanding and predicting demand-side behaviors in transportation, energy, and other infrastructure systems (Roorda et al., 2008, Pawlak et al., 2021).

While recent advancements have focused on intra-household interactions, such as the coordination of schedules between household members (as seen in the work of Rezvany et al., 2023), there is a growing need to extend these models to capture interactions beyond the household. As modern urban environments become more interconnected, individuals coordinate with a wide range of people—friends, colleagues, and mobility-sharing groups—when planning their day. At the same time, the increasing richness of modern travel datasets further motivates the development of models capable of capturing such interactions. Given this context, several challenges arise in extending activity-based modeling frameworks. First, how can we effectively capture the social interactions that influence individuals' daily decisions, particularly those that go beyond household settings? Second, how can the model be made flexible enough to accommodate various types of interactions and choice dimensions while ensuring scalability for real-world applications?

While existing frameworks, such as that of Rezvany et al. (2023), have made significant progress in capturing joint activity participation within households, they only partially account for more flexible forms of coordination—for example, when agents start a shared activity at different times or choose different transport modes after meeting. To address this limitation, our work extends and generalizes their framework to represent a broader range of scheduling behaviors. The proposed model, formulated as a Mixed Integer Linear Problem (MILP) with a distinct and more efficient structure, enables richer behavioral representations and improved computational performance. Beyond households, these modeling capabilities are also applicable to other forms of social interaction—among friends, colleagues, or participants in shared mobility systems—where coordination extends beyond traditional family boundaries (Cirillo and Axhausen, 2006).

Several works formulate the optimization problem underlying ABMs as an MILP (Recker, 1995, Pougala et al., 2022). However, limited attention has been devoted to improving the structure and compactness of these formulations. Our model introduces a compact version of the model by Rezvany et al. (2023), which can be easily extended with new features. By using a graph-based approach, we also open the door to the diverse techniques developed over decades to improve the solving of graph optimization problems, such as dynamic programming.

We propose a versatile and easy-to-apply model for casting the collective decision of a (small) group of people who schedule their activities over a time period so as to maximize their global utility. Methodologically, we formulate the problem as a *multi-agent scheduling problem*. To solve it, we adopt a graph-based representation inspired by the approach of Gaul et al. (2022) to solve the Dial-a-Ride Problem. The multi-agent scheduling problem is hence reduced to a minimum-cost flow problem in a labeled directed graph, with some additional constraints. We solve the resulting MILP using a commercial mixed-integer linear solver.

This paper makes three main contributions. First, we formulate the multi-agent scheduling problem as a graph-based mixed-integer linear program that compactly represents activities, travel, and coordination decisions. Second, the proposed formulation generalizes existing household-based models by allowing flexible synchronization patterns, including staggered participation and heterogeneous transport choices after the shared activities. Third, numerical experiments demonstrate that the resulting model achieves computational improvements over prior formulations while enabling richer behavioral representations. By combining the behavioral realism of ABMs with the structural clarity and scalability of graph-based optimization, our work contributes a flexible and extensible tool for modeling multi-agent activity scheduling.

The remainder of the paper is structured as follows. In Section 2, we first review the existing activity-based modeling literature and identify the gaps in current research. We then position our work within the operations research literature on vehicle routing problems (VRP), emphasizing optimization formulations that are formally close to ours. Section 3 introduces the multi-agent scheduling problem, providing the problem description and relevant notation. In Section 4, we present the graph-based model and the solution methodology, detailing the graph-based optimization approach and the MILP. Finally, Section 5 presents the numerical experiments, comparing our model with Rezvany et al. (2023) on real-world data to demonstrate its flexibility and scalability in capturing a wider range of social interactions.

2 Literature review

2.1 Activity-based models

ABMs have become a central tool in transportation research due to their ability to simulate individuals' activities and travel patterns over time and space. These models are grounded on the behavioral assumption, conceptualized by Chapin (1974), that travel demand is not an isolated phenomenon but is driven by participation in various activities, which are distributed spatially and temporally. Hence, ABMs focus on individuals and their choice of activities, while taking into account their spatial and temporal constraints.

The historical development of ABMs dates back to works such as the model by Adler and Ben-Akiva (1979) and other early models reported in the study by Axhausen and Gärling (1992). These works emphasize the importance of understanding activity participation and scheduling for transport demand modeling. They laid the groundwork for the activity-based approach, which has since evolved to include increasingly complex representations of human behavior.

Despite their advantages, traditional ABMs face several limitations. A key challenge is that many existing models focus on individual decision-making without considering the interactions between agents (Bhat, 2005, Habib and Hui, 2017). This simplification assumes that individuals make decisions in isolation, disregarding the fact that activities and travel plans are often coordinated among family members or social groups. For example, when household members are treated as independent agents, the model fails to capture the coordination required for joint travel, childcare, and shared household responsibilities (Srinivasan and Athuru, 2005).

Early household-level models address intra-household interactions, such as car allocation, joint activities, and travel arrangements. Notable works include Vovsha et al. (2003) and Gupta et al. (2014), who introduce models that consider the coordination of activities and travel within the household. These models still fall short in capturing the complexity of group decision-making, particularly when considering heterogeneous time constraints and shared activity participation.

Another limitation arises from the sequential scheduling approach used in many ABMs

(Vovsha et al., 2005). These models treat scheduling as a series of discrete steps, where decisions regarding one activity or trip are made before the next. While convenient, this approach fails to capture the simultaneous interdependencies between different scheduling decisions, such as the time of day an activity occurs and the transport mode used. A simultaneous approach, in which all decision dimensions are considered jointly, makes it easier to model trade-offs compared to a model where a sequencing order is assumed.

To address the limitations of sequential ABMs, ABMs have been introduced, where all decision dimensions—activity choice, timing, location, and mode—are considered jointly (Ettema et al., 2007). Examples include the Multiple Discrete-Continuous Extreme Value models (MDCEV) applied to activity scheduling (e.g., Lai et al., 2019, Palma et al., 2021) and the Household Activity Pattern Problem (HAPP) formulated as an MILP (Recker, 1995, Gan and Recker, 2013). Although Lai et al. (2019) incorporate intra-household interactions and all decision dimensions, their approach does not solve the activity scheduling problem using optimization tools. Instead, once the behavioral parameters are estimated, daily schedules are generated through probabilistic simulation based on the MDCEV choice structure, without computing an optimal joint schedule for the household.

The HAPP models have been widely used to model activity scheduling as MILPs, drawing from vehicle routing literature (Recker, 1995). However, HAPP models often face challenges due to their prescriptive nature and difficulties in calibration. These models typically assume a fixed set of activities and do not account for the variability in activity participation or duration. Recent variants, such as those by Xu et al. (2018), allow for some flexibility in activity participation and duration but still use a sequential simulation approach, which limits their ability to capture simultaneous interdependencies.

Building on previous optimization-based approaches, Pougala et al. (2022) propose an activity-scheduling framework that integrates multiple choice dimensions—activity participation, location, start time, duration, and mode—into a single optimization problem, capturing the complex trade-offs between scheduling decisions for multiple activities. The model represents all time-related choices as continuous variables, with a modular utility formulation that allows preferences to be specified for each individual activity, and is implemented as an MILP that can be extended with custom constraints. While Pougala et al. (2022) demonstrated the framework on small samples, Rezvany et al. (2023) extended it to household-level modeling, allowing the simultaneous scheduling of activities for all household members and accounting for intra-household coordination. This framework, referred to as OASIS (Optimization-based Activity Scheduling with Integrated Simultaneous choice dimensions), represents the most recent advance in flexible and integrated activity-based modeling, combining multiple decision dimensions while maintaining a behavioral basis (Pougala et al., 2023). Additionally, Manser et al. (2024) demonstrated the operational relevance of this approach through an application of the Pougala et al. model to the Swiss Railways, showcasing its practical applicability in transportation scheduling.

Existing frameworks, including the OASIS model for households (Rezvany et al., 2023), have advanced the modeling of coordinated activity-travel behaviors but remain limited in the diversity of interactions they can represent. In particular, they struggle to capture situations where individuals coordinate flexibly in time, space, or mode choice—for instance, when participants engage in different activities at the same location, join or leave a shared activity at different times, or switch transport modes during a trip.

To overcome these limitations, our formulation generalizes and extends the OASIS framework to encompass a much wider range of coordination patterns. It can represent: (a) agents traveling together to perform different activities at a shared location; (b) staggered participation in joint activities; (c) escorting behavior, where one agent accompanies another before continuing elsewhere; (d) synchronized or independent execution of identical activity sequences; (e) constraints requiring the presence of specific agents (e.g., escorting, caregiving, or shared vehicle use); and (f) multimodal trip chaining, such as combining cycling with public transport.

Conceptually, the key innovation lies in the model's structure: whereas OASIS links activities directly to their subsequent trips, our approach models activities and trips separately, with activities represented as vertices and trips as arcs in a graph. This graph-based formulation enables greater modeling flexibility, allowing for richer and more realistic coordination dynamics among agents.

2.2 Related scheduling and routing problems

While optimization-based approaches, such as HAPP and OASIS models, enable simultaneous decision-making for activities, locations, and modes, they involve combinatorial decisions. This makes the models computationally expensive, especially when scaling up to larger numbers of agents or more complex scheduling scenarios. One of the key challenges in optimization-based ABMs is hence ensuring flexibility while maintaining scalability.

A promising research direction lies in the vehicle-routing literature, where highly combinatorial problems are routinely solved with efficient algorithms. In particular, the Vehicle Routing Problem with synchronization (VRPs) and the Dial-a-Ride Problem (DARP) share fundamental structural properties with ABMs: agents perform both individual and shared tasks, must respect spatial and temporal constraints, and coordinate with others for joint activities. Moreover, the number of participants in each shared activity remains small (e.g., passengers in a vehicle, for DARP, or members of a working team, for VRPs), which makes these formulations aligned with multi-agent ABMs.

In the DARP literature, Gaul et al. (2022) show that event-based formulations outperform traditional models by representing each request through its pickup and drop-off events rather than through spatial nodes. This event structure encodes precedence, capacity, and pairing constraints directly in the graph, removing the need for explicit constraints and yielding remarkable performance on standard DARP benchmarks, even though it introduces a larger number of variables.

Our model adopts a similar event-based philosophy for graph construction. Each vertex in our graph corresponds to a specific location and a subset of agents who may jointly perform an activity there. This design mirrors the event-based structure in DARP, where each vehicle state is tied to a passenger subset. The number of vertices remains manageable because the subset size is bounded (by vehicle capacity in the DARP, or by household size in our case). However, unlike in the graph of Gaul et al. (2022), our vertices do not represent instantaneous events with a start time to be determined. They represent activities for which both the start time and the duration must be scheduled. As a consequence, the graph structure proposed by Gaul et al. (2022) cannot be used directly in our setting, and a new graph is required to incorporate the additional complexities of our problem.

Once the schedule of each agent has been modeled as a path in a graph, the problem becomes closely related to the vehicle routing problem with synchronization. The taxonomy proposed by Drexl (2012) distinguishes between several types of synchronization, among which exact operation synchronization is the most relevant to coordinated activity scheduling. It requires that multiple agents perform actions simultaneously, such as jointly participating in an activity. In the context of homecare—the most common application of exact operation synchronization—several nurses with different skills must coordinate to jointly deliver a care task to a patient at the same time. The formulation by Bredström and Rönnqvist (2008) exemplifies this concept, developed for homecare staff scheduling and forest operations, enforcing both pairwise synchronization and temporal precedence between visits. Their MILP formulation is mathematically the closest to ours once the graph representation is established, differing mainly in the objective function and in the fact that their model requires every vertex to be visited—a constraint later relaxed in their heuristic solution. However, it is important to note that our formulation does not require this constraint, which is a key distinction between ABMs and VRPs.

Table 1 summarizes the main characteristics of the three key interrelated works: the ABM of Rezvany et al. (2023), which addresses the same multi-agent scheduling problem (MASP) at the household level; the event-based DARP formulation of Gaul et al. (2022), which inspired our vertex construction; and the VRP with synchronization of Bredström and Rönnqvist (2008), whose MILP structure most closely matches ours. Together, these works form the conceptual foundation for the proposed graph-based approach to the MASP.

In this comparison, direct formulations define decision variables directly over activities or trips, without requiring any preprocessing or abstract graph representation. In contrast, event-based formulations introduce an intermediate graph where each vertex represents an event and a subset of present agents. Our approach follows this latter structure while preserving the behavioral richness of activity-based models, i.e., preference of agents in the objective function, modeling of shared trips and activities, choice of start, duration, location, and mode.

Article	**Rezvany et al. (2023)**	**Gaul et al. (2022)**	**Bredström & Rönnqvist (2008)**	**This work**
Problem	MASP	DARP	VRPs	MASP
Formulation	Direct	Event-based	Direct	Event-based
Solving method	MILP solver	MILP solver	MILP solver + heuristic	MILP solver
Preferences of agents	Yes	No	No	Yes
Shared trips	By car only.	Yes	No	Yes
Shared activities	Yes	No	Yes	Yes
Schedule	All dimensions (start, duration, and location)	Only order of events	Only order of tasks	All dimensions (start, duration, and location)
Choice of mode	Yes	No	No	Yes
Coordination flexibility	Limited (shared activity implies shared trip)	No	High (independent trip and activity sharing)	High (independent trip and activity sharing, multimodal, staggered)

Table 1: Comparison of the modeling features of relevant articles.

3 The multi-agent scheduling problem

Activity-based models typically simulate a large number of groups of individuals. In this work, we assume that group decisions are independent across groups and restrict our analysis to the interactions occurring within each group. We are hence interested in the detailed scheduling problem of one group of individuals over a period of time. Our model aims to determine simultaneously which activities the individuals perform, their location, their schedule, and the transport mode used to travel from one activity to another. We make the assumption that the group's objective is to maximize a joint utility function, which depends on the utility function of all individual members.

To model this problem, we introduce the concept of *agents*, which encompasses both actual individuals and non-human resources such as private vehicles, EV chargers, or household appliances (e.g., laundry machines). A resource is defined as any entity whose availability or schedule constrains that of the individuals. This generalization makes it possible to represent, for instance, the joint scheduling of several household members sharing one or more private vehicles. In such cases, the agents of the system comprise both the individuals and their vehicles, whose schedules must be coordinated.

In this section, we introduce the multi-agent scheduling problem (Section 3.1). As already stated in the introduction, this problem captures many concrete situations where a group of people schedules its activities collectively. Section 3.2 explains why auxiliary activities are needed in the model and how they are defined. A concrete example is given in Section 3.3, where the high versatility of this problem is illustrated and emphasized.

3.1 Formalization

We denote by N the set of agents who want to schedule their activities and trips over a continuous time period $[0, T]$ while maximizing their joint utility. We detail now the characteristics of the agents, the activities, the trips, and the utilities. All notation is provided in Table 2.

3.1.1 Agents

We define an agent as any entity within the group that:

- has a schedule, i.e., a sequence of activities or states evolving over time,
- consumes or provides resources during these activities,
- imposes constraints on the schedules of other agents due to limited availability,
- requires coordination with other agents in the group to ensure overall feasibility.

In our framework, the notion of agent is flexible and can encompass both individual and non-human resources. For instance, individuals typically participate directly in decision-making, and their utilities contribute to the objective function with positive weights. By contrast, non-human resources—such as vehicles, EV chargers, parking spaces, or teleworking offices—can be considered as coordinating agents that enable or constrain activity schedules; their utilities may therefore be assigned a weight of zero.

For each agent n, we define a set of possible activities A_n, a home location $\text{home}(n)$, and a financial budget B_n over the time period. The budget constraint is here modeled at the individual level; however, it could equivalently be formulated at the group or household level without affecting the generality of the framework.

3.1.2 Activities

The set of all activities is denoted by $A := \bigcup_n A_n$. Each activity can be performed only at some specific *locations* and the set of all possible locations for an activity a is denoted by L_a. For instance, the activity of grocery shopping can be performed at any supermarket located in the area where the agent lives or works.

Each activity a comes with a *participation limit* C_a, which is the maximal number of agents of the group who can perform simultaneously the activity. This is particularly useful to model activities related to private vehicles, with a limited capacity.

Each activity a is associated with a collection N_a of *required agent configurations*. To be performed, at least one of these configurations must be selected, meaning that all agents included in the chosen configuration jointly participate in the activity—at the same location and time, while other agents may possibly perform the activity with them. This formulation is particularly relevant for modeling trips by private vehicle, where both the vehicle itself and a driver must be available. For instance, in a household with two adults holding a driving licence, an activity a related to driving could have $N_a = \{\{\texttt{car}, \texttt{adult 1}\}, \{\texttt{car}, \texttt{adult 2}\}\}$. It also captures situations such as household dinners, where at least one adult needs to be present to prepare the meal for young children.

If an activity a is performed at a location $\ell \in L_a$, then it must be during *opening hours* $[\gamma^-_{a\ell}, \gamma^+_{a\ell}]$, and its duration must last at least $\tau^{\min}_a$ and at most $\tau^{\max}_a$. The majority of out-of-home activities have restrictive opening hours (e.g., shopping, restaurants, leisure activities). Each agent n that can perform an activity a has a *preferred starting time* $x^\star_{an}$ and a *preferred duration* $\tau^\star_{an}$. For instance, activities like jogging or sleeping don't have strict schedules, but it is more convenient to perform them during certain times of the day. Performing an activity a at location ℓ can also have a *cost* $c_{a\ell}$ to be paid by each agent performing the activity (for example, the cost of a leisure activity). We define K groups G_k of *mandatory activities*, linked with an integer $p_k \leqslant |G_k|$. Among this group, at least p_k activities have to be performed. Note that if G_k contains only one activity a and $p_k = 1$, then activity a has to be performed by an agent of the group. This feature is particularly useful for modeling household chores that need to be done periodically, such as cleaning tasks.

Each agent n can choose which activities she performs among her set A_n of activities and the starting time and the duration of these activities. Temporal consistency is enforced: an agent can start a new activity only after completing the previous one and traveling to the new location. Each activity is assumed to be performed at most once within the time period. Formally, the activity set A is defined such that every element $a \in \mathcal{A}$ corresponds to a unique occurrence of an activity. When an activity type recurs in reality—such as work performed in both the morning and the afternoon, each occurrence is modeled as a distinct element of A. This ensures that the scheduling problem remains well-defined, while still allowing the model to capture repeated behaviors within the same horizon.

The schedule of every agent n must cover the entire time period, starting with an activity $\texttt{dawn} \in A_n$ and ending with an activity $\texttt{dusk} \in A_n$, both of them performed at location home(n), the home of agent n.

3.1.3 Trips

To travel between two locations k and ℓ, agent n has to choose one of the *available transport modes* $M^{k\ell}_n$. A trip from k to ℓ and a transport mode m have an associated *cost* $\rho_{k\ell m}$ and *duration* $d_{k\ell m}$.

We also assume that congestion is exogenous—that is, travel costs and travel times may vary by time of day (e.g., during peak hours), but they are not affected by the group's own travel decisions.

3.1.4 Utility

The utility U_n of each agent n is defined following Pougala et al. (2022), and consists of two deterministic and two random components. All the elements entering this definition, apart from the decision variables themselves, are exogenous: they do not depend on the agents' choices, must be precomputed, and are provided as inputs to the model. In practice, the utility functions are calibrated from behavioral data, and the parameters are estimated using methods such as discrete choice modeling (e.g., Greeven et al., 2005, Guarda and Qian, 2024).

For clarity, we introduce the following decision variables:

- x_{an}: starting time chosen by agent n for activity a,
- τ_{an}: duration chosen by agent n for activity a,
- w_{an}: binary variable equal to 1 if agent n performs activity a, and 0 otherwise,
- $j_{ann'}$: binary variable equal to 1 if activity a is performed by agent n jointly with another agent n',
- $z_{\ell\ell' mn}$: binary variable equal to 1 if agent n makes the trip from ℓ to ℓ' using mode m, and 0 otherwise,
- $j_{\ell\ell' mnn'}$: binary variable equal to 1 if the trip of agent n from ℓ to ℓ' using mode m is shared with another agent n'.

Based on these variables, the utility of agent n is composed of:

- The sum of the utility $U^n_{a\ell}$ for each activity a at location $\ell \in L_a$ performed by agent n.
- The sum of the random terms ξ_{an} over all the activities a performed by agent n.
- The sum of the utility $U^n_{\ell\ell' m}$ for each trip (ℓ, ℓ') performed by agent n with transport mode m.
- The sum of the random terms $\xi_{\ell\ell' mn}$ for each trip (ℓ, ℓ') performed by agent n with a mode m.

Hence, it is written as:

$$U^n = \sum_{a \in A_n} \sum_{\ell \in L_a} w_{an} \left(U^n_{a\ell} + \xi_{an} \right) + \sum_{\ell, \ell', m} z_{\ell\ell' mn} \left(U^n_{\ell\ell' m} + \xi_{\ell\ell' mn} \right).$$

Let us now describe these components in detail. Each agent n derives a utility $U^n_{a\ell}$ from performing activity a at location ℓ, written as:

$$
\begin{aligned}
U^n_{a\ell} \coloneqq\ & r_{a\ell n} + \sum_{n' \in N} r^{\text{joint}}_{ann'}\, j_{ann'} + \theta_c\, c_{a\ell} \\
& + \theta_{x\,\text{early}}\big(x^*_{an} - x_{an}\big)^+ + \theta_{x\,\text{late}}\big(x_{an} - x^*_{an}\big)^+ \\
& + \theta_{\tau\,\text{short}}\big(\tau^*_{an} - \tau_{an}\big)^+ + \theta_{\tau\,\text{long}}\big(\tau_{an} - \tau^*_{an}\big)^+.
\end{aligned}
\tag{1}
$$

This utility is composed of:

- An exogenous reward $r_{a\ell n}$, received for performing the activity a at location ℓ.
- An exogenous joint reward $r^{\text{joint}}_{ann'}$ if the agent performs activity a together with another agent n'.
- A disutility $\theta_c c_{a\ell}$ proportional to the cost $c_{a\ell}$ of activity a at the chosen location ℓ, where θ_c is exogenous and expected to be negative.
- Two disutilities for the deviations from the desired schedule. To ease the definition of these penalties, we introduce x_{an} and τ_{an} the starting time and duration chosen by agent n for activity a.
  - ⋆ A disutility $\theta_{x\,\text{early}}(x^*_{an} - x_{an})^+ + \theta_{x\,\text{late}}(x_{an} - x^*_{an})^+$ proportional to the deviation from the preferred starting time x^*_{an}, with exogenous negative weights for starting the activity too early, $\theta_{x\,\text{early}}$, and too late, $\theta_{x\,\text{late}}$. Here, $(\cdot)^+$ denotes the positive part operator, defined as $(y)^+ = \max(y, 0)$.
  - ⋆ A disutility $\theta_{\tau\,\text{short}}(\tau^*_{an} - \tau_{an})^+ + \theta_{\tau\,\text{long}}(\tau_{an} - \tau^*_{an})^+$ proportional to the deviation from the preferred duration τ^*_{an}, with exogenous negative weights for a shorter duration, $\theta_{\tau\,\text{short}}$, and a longer duration, $\theta_{\tau\,\text{long}}$.

Similarly, each agent n derives a utility $U^n_{\ell\ell' m}$ from the trip between locations ℓ and ℓ' using transport mode m, written as:

$$
U^n_{\ell\ell' m} \coloneqq r_{\ell\ell' mn} + \sum_{n' \in N} \big(r^{\text{joint}}_{\ell\ell' mnn'}\, j_{\ell\ell' mnn'}\big) + \theta_{ttm}\, d_{\ell\ell' m} + \theta_{tc}\, \rho_{\ell\ell' m}. \tag{2}
$$

This utility is the sum of:

- A reward $r_{\ell\ell' mn}$, capturing the utility derived by agent n when using transport mode m to travel from location ℓ to location ℓ'. It reflects all factors other than cost and time, such as comfort, reliability, or convenience, and is considered exogenous.

- An exogenous joint reward $r^{\text{joint}}_{\ell\ell' mnn'}$, received if the agent shares the trip with another agent n'. This reward term captures the social benefits of shared travel, as agents may gain additional utility from traveling together, as highlighted by Vovsha et al. (2003).

- A disutility $\theta_{ttm} d_{\ell\ell' m}$ proportional to the the travel time $d_{\ell\ell' m}$, with a negative exogenous weight θ_{ttm} depending on the mode m used for the trip.

- A disutility $\theta_{tc}\rho_{\ell\ell' m}$ proportional to the travel cost $\rho_{\ell\ell' m}$, with a negative exogenous weight θ_{tc}.

Note that the rewards for joint activities and trips can be zero or even negative if the agent n doesn't want to share moment with other individuals.

These utility functions can be calibrated from behavioral data, where observed travel choices and preferences are used to estimate the parameters. For example, methods such as discrete choice modeling have been applied to calibrate travel behavior in the context of ride-sharing (see, e.g., Bhat and Sardesai, 2006; Olaru et al., 2012).

For each agent n and activity a in A_n, the random term ξ_{an} has a known distribution and captures the errors in modeling the preferences of agent n regarding activity a. Identically, for each agent n and trip (ℓ, ℓ') performed with a mode m, the random term $\xi_{\ell\ell' mn}$ has a known distribution and captures the errors in modeling the preferences of agent n regarding trip (ℓ, ℓ') with mode m.

We assume that the group of individuals selects the schedule that maximizes its overall utility, defined as an aggregation of the individuals' utilities. To make this assumption, individuals must act in the interests of the group, so there is no competition between them. This may be the case for a household, a group of friends, or co-workers. Although there is no consensus in the literature on how to define such an aggregation function (see, e.g., De Palma et al., 2014), several approaches have been discussed: an additive rule based on the (possibly weighted) sum of the utilities of each member, an autocratic rule privileging the utility of the "strongest" member, or an egalitarian rule privileging the "weakest" member (Shilov et al., 2025, Vo et al., 2020, Kurata and Nakamura, 2025). In our framework, we adopt the additive formulation, but the other formulations can be considered as well. We associate with each agent $n \in N$ a weight $\omega_n \geqslant 0$, which represents the relative importance of agent n in the aggregated utility. The total utility U is then defined as the weighted sum of the agents' utilities:

$$U = \sum_{n \in N} \omega_n U^n. \tag{3}$$

This formulation also allows us to assign a weight of zero to non-human resources, such as private vehicles, ensuring that they are represented in the model without directly contributing to the objective.

Notation	Description
N	Set of agents
A_n	Set of possible activities of agent n
A	Set of all activities
L_a	Set of locations where activity $a \in A$ can be performed
$c_{a\ell}$	Cost for an agent to perform activity a at location $\ell \in L_a$
$\gamma^-_{a\ell}$	Opening time for activity a at location ℓ
$\gamma^+_{a\ell}$	Closing time for activity a at location ℓ
τ^{min}_a	Minimum duration for activity a
τ^{max}_a	Maximum duration for activity a
$x^\star_{an}$	Time at which agent n would prefer to start activity a
$\tau^\star_{an}$	Duration agent n would like to perform activity a
$G_k \subseteq A$	Subset of activities of A that need to be performed at least p_k times
C_a	Maximum number of agents that can perform activity a
$N_a \subseteq 2^N$	Collection of subsets of agents required to jointly perform activity a
$M^{k\ell}_n$	Set of available transport modes for agent n to travel between locations k and ℓ
$\rho_{k\ell m}$	Cost of traveling from k to ℓ using transport mode m
$d_{k\ell m}$	Duration of traveling from k to ℓ using transport mode m
T	Time period
B_n	Daily budget of an agent n
ω_n	Weight of an agent n in the total utility
$r_{a\ell n}$	Reward for agent n performing activity a
$r^{\text{joint}}_{ann'}$	Additional reward if agent n performs activity a with another agent n'
θ_c	Penalty for the activity cost
$\theta_{x\text{ early}}, \theta_{x\text{ late}}$	Penalties for the deviation from the preferred starting time
$\theta_{\tau\text{ short}}, \theta_{\tau\text{ long}}$	Penalties for the deviation from the preferred duration
θ_{ttm}	Penalty for travel time with mode m
θ_{tc}	Penalty for travel cost
$r_{\ell\ell' mn}$	Reward for agent n traveling with mode m from ℓ to ℓ'
$r^{\text{joint}}_{\ell\ell' mnn'}$	Additional reward when agent n travels with agent n' on trip (ℓ, ℓ', m)
ξ_{an}	Random term with known distribution
$\xi_{\ell\ell' mn}$	Random term with known distribution

Table 2: Notation.

3.2 Modeling of special and auxiliary activities

While many activities are directly interpretable as actions performed by individuals (e.g., working, shopping, or eating), certain situations require the introduction of *auxiliary activities*. These are activities that do not correspond to a literal action but that influence the schedule of the group members. An auxiliary activity is any activity that is either (i) performed by a non-human resource (for example, an activity "parking" can be defined for a private vehicle), or (ii) indirectly affects the feasible schedules of human agents without representing a real action (such as an activity "escorting" that requires an adult to bring a child to school).

It is assumed that the modeler has sufficient contextual knowledge to define these activities and their main attributes, including location, duration, and required agents. For auxiliary activities related to private vehicles, the *Vehicle-based Auxiliary Activity Generation Algorithm*, described in Appendix 6, automatically generates a complete set of auxiliary activities based on the real activities of the group members who use the vehicle. This guarantees that private vehicles can be modeled, and we can assume their sets of activities to be well defined. Nevertheless, the algorithm is intentionally general and may produce more auxiliary activities than needed in practice. Consequently, the final specification relies on the modeler's understanding of household behavior to define a realistic and parsimonious activity set.

3.3 Concrete example

We provide a detailed example of carpooling between two colleagues, illustrating how to define sets of activities, locations, and transport modes to model the schedule of a private vehicle. The data and parameters are presented here in a straightforward manner, but in Section 4.3, we explore how these choices can be interpreted, demonstrating that they enable the model to closely align with reality.

Consider a simple instance involving three agents: Alice, Bob, and Alice's car (represented by an agent Car). Alice and Bob are colleagues who are scheduling their day together to maximize their total utility. Alice and Bob can work at the office, together or separately, and Bob also has the option of playing tennis in the city area. Transportation options include public transport (PT), which is available to both Alice and Bob, and Alice's car, which only Alice is allowed to drive. However, Bob can travel by car when he shares his trip with Alice. Each agent is assigned a time budget T of 24 hours.

3.3.1 Agents

In this example, the agents are Alice, Bob, and Car. Alice and Bob are individuals with a financial budget $B_n = 10$ and home locations `homeAlice` and `homeBob`, respectively. Car is a non-human resource with home location `homeAlice` and an infinite financial budget.

Activities of Car are supposed to be generated in the preprocessing by the modeler. Hence, their respective sets of activities are:

$$A_{\text{Alice}} = \{\texttt{dawn, dusk, work,} \\ \texttt{departure from homeAlice,} \\ \texttt{arrival to office,} \\ \texttt{departure from office,} \\ \texttt{stop to drop, arrival to homeAlice}\}$$

$$A_{\text{Bob}} = \{\texttt{dawn, dusk, work, tennis,} \\ \texttt{departure from office, stop to drop}\}$$

$$A_{\text{Car}} = \{\texttt{dawn, dusk, departure from homeAlice,} \\ \texttt{arrival to office, parking,} \\ \texttt{departure from office, stop to drop,} \\ \texttt{arrival to homeAlice}\}$$

3.3.2 Activities

The activities are grouped by location as follows:

- At homeAlice, the activities are dawn, dusk, departure from homeAlice, and arrival to homeAlice.
- At homeBob, the activities are dawn and dusk.
- At the office, the activities are work, arrival to office, parking, and departure from office.
- At the city, the activities are tennis and stop to drop.

To simplify the example, we assume non-restrictive time-windows on the activities. For all activities and locations, the opening hours span the entire day from time 0 to 24, and both the minimum and maximum durations are set to 0 and 24, respectively.

We assume that the cost of activity tennis is $c_{\texttt{tennis city}} = 3$, and $c_{a\ell} = 0$ for every other activity a and location ℓ. For each activity a, the set N_a of required agents is empty by default, except for activities related to driving, like `departure from homeAlice`, `arrival to office`, `departure from office`, `stop to drop`, `arrival to homeAlice`, for which $N_a = \{\{\text{Alice}, \text{Car}\}\}$. This enforces that both Alice and Car are required to perform these activities. To represent the limited capacity of the car, the participation limit for these activities, is set to 5, i.e., $C_a = 5$. All other activities have an unbounded participation limit, i.e., $C_a = \infty$ for all $a \in A$ where the set N_a is empty.

3.3.3 Trips

The available transport modes are the `car`, the public transport (`PT`), and the mode `noTrip`, which is available to each agent n to go from a location ℓ to the same location, i.e., $M_n^{\ell\ell} = \{\texttt{noTrip}\}$ for all ℓ. The set of available transport modes for each agent between different locations are described in Table 3.

Origin Destination	homeAlice office	office city	office homeBob	city homeBob	city homeAlice
Alice	{PT, car}	{PT, car}	∅	∅	{PT, car}
Bob	∅	{PT, car}	{PT}	{PT}	∅
Car	{car}	{car}	∅	∅	{car}

Table 3: Available transport modes for each agent.

The travel times $d_{\ell\ell'\text{PT}}$ and $d_{\ell\ell'\text{car}}$ between two locations ℓ and ℓ' are defined in Figure 1. By definition, the distance between a location and itself is zero, and travel times are symmetrical. For any couple (ℓ, ℓ') of locations, we also define the costs $\rho_{\ell\ell'\text{PT}} = d_{\ell\ell'\text{PT}}$ and $\rho_{\ell\ell'\text{car}} = 2\ d_{\ell\ell'\text{PT}}$.

Figure 1: Travel time map for transport mode `car` (left) and `PT` (right).

3.3.4 Utility

The preferred starting time and duration of each activity are defined in Table 4, providing an overall preferred schedule for Alice and Bob. Alice prefers to leave home at 8:00 and starts working at 9:30, while Bob has a slightly different schedule, with his preferred start time for work at 8:00. Bob also prefers to play tennis between 17:00 and 19:00 in the city. Other activities, such as stopping to drop off, arriving at work, and returning home, are similarly defined to ensure a cohesive daily routine for both agents.

Name	**Pref. starting time**		**Pref. duration**	
	Alice	**Bob**	**Alice**	**Bob**
`dawn`	0:00	0:00	7h	7h
`dusk`	20:00	21:00	3h	3h
`work`	9:30	8:00	8h	8h
`tennis`	-	17:00	-	2h
`departure from homeAlice`	8:00	-	0h	-
`arrival to office`	9:30	-	0h	-
`departure from office`	17:00	16:00	0h	0h
`stop to drop`	17:00	17:00	0h	0h
`arrival to homeAlice`	19:00	-	0h	-
`parking`	-	-	-	-

Table 4: Preferred starting time and duration for each activity and agent.

If the example was to be fully specified, we would need to define the parameters in the utility function, including the weights for activity duration deviations ($\theta_{\tau\,\text{short}}$, $\theta_{\tau\,\text{long}}$), starting time deviations ($\theta_{x\,\text{early}}$, $\theta_{x\,\text{late}}$), and travel time ($\theta_{\mathtt{ttm}}$). Additionally, the rewards for each activity and trip would need to be specified. These parameters, typically estimated from data, reflect the relative importance of each factor in the agents' utility functions.

The weight $\omega_{\mathtt{car}}$ is 0, indicating that Alice's car does not derive any utility. The weights $\omega_{\mathtt{Alice}}$ and $\omega_{\mathtt{Bob}}$, however, are equal and non-zero.

4 Graph-based model

In this section, we model the multi-agent scheduling problem as a path problem in a directed graph, and then show how to write this flow problem as an MILP.

A path in this graph captures the choice of activities that an agent performs, their locations, and the order in which the activities are performed. The scheduling of the activities for each agent will be modeled by additional decision variables on each vertex of the graph. This approach has been inspired by the event-based modeling of the DARP introduced by Gaul et al. (2022). It has the advantage of encoding some complicating constraints in the graph structure.

4.1 Constrained minimum-cost flow formulation

The graph, denoted by $G = (V, E)$, has labeled arcs, and is defined as follows. For each activity a and location $\ell \in L_a$, we introduce a vertex v as a triple $v = (a, \ell, S)$, for all subsets $S \subseteq N$ satisfying the following conditions:

- S is a subset of agents that can perform a, i.e., $S \subseteq \{n \in N \colon a \in A_n\}$,

- the participation limit of activity a is not exceeded, i.e., $|S| \leqslant C_a$, and
- there is a subset $X \subseteq S$ of required agents in S, i.e., $X \in N_a$.

A vertex $v = (a, \ell, S)$ has an obvious meaning: a subset S of agents performing together activity a at location ℓ. It is also associated with a utility u_v. This vertex utility corresponds to the part of the activity utility defined in Equation (1) that is independent of the scheduling variables (start time and duration), and this, for all the agents in S. Formally, it is defined as

$$u_v \coloneqq \sum_{n \in S} \omega_n \Big(r_{a\ell n} + \sum_{n' \in S} r^{\text{joint}}_{ann'} + \theta_c \, c_{a\ell} \Big).$$

We introduce a specific notation for the vertices related to the activities `dawn` and `dusk`. For each agent n, the activity `dawn` represents the first activity of the considered time period, performed alone at the agent's home location $\text{home}(n)$. It is therefore defined as

$$\texttt{dawn}(n) \coloneqq (\texttt{dawn}, \text{home}(n), \{n\}).$$

Symmetrically, the activity `dusk` is defined as

$$\texttt{dusk}(n) \coloneqq (\texttt{dusk}, \text{home}(n), \{n\}).$$

Let $v = (a, \ell, S)$ and $v' = (a', \ell', S')$ be two vertices such that $S \cap S' \neq \varnothing$, i.e., there is at least one agent that can perform activity a at location ℓ and activity a' at location ℓ'. For every mode $m \in \bigcap_{n \in S \cap S'} M_n^{\ell\ell'}$, i.e., a mode that can be used by the agents common to S and S' to go from ℓ to ℓ', we introduce an arc $e = (v, v')$ with a label $m_e \coloneqq m$. This arc e is associated with a cost equal to $\rho_e \coloneqq \rho_{\ell\ell' m}$ and a travel time (or duration) $d_e \coloneqq d_{\ell\ell' m}$, which correspond to the cost and time of traveling from ℓ to ℓ' using mode m for one agent. It is also associated with a utility u_e, containing the term of the trip utility defined in Equation (2) that does not dependent on the other agents (the joint reward is counted separately). It is written

$$u_e \coloneqq \theta_{ttm} d_{\ell\ell' m} + \theta_{tc} \rho_{\ell\ell' m}.$$

Note that parallel arcs are allowed, i.e., there can be several arcs between two vertices, each one associated with a different mode.

The problem becomes the following: for each agent n, compute a directed $\texttt{dawn}(n)$-$\texttt{dusk}(n)$ path in G, and a starting time x_v and a duration τ_v for each visited vertex v, under "combinatorial" constraints, "time-consistency" constraints, and a "budget" constraint, so as to maximize the total utility.

The combinatorial constraints are the following:

- the path of an agent n only visits vertices (a, ℓ, S) with $n \in S$.

- if the path of an agent n visits a vertex (a, ℓ, S), then the paths of all agents in S must visit that vertex.

- for each activity a and subset S of agents, if a vertex (a, ℓ, S) is visited by a path, no vertex (a, ℓ', S) with $\ell \neq \ell'$ can be visited. This ensures that the agents in S choose a single location ℓ to perform activity a.

- for each group G_k of activities, at least p_k vertices (a, ℓ, S) such that $a \in G_k$ must be visited by some agent.

The time-consistency constraints are of three types:

- if a vertex $v = (a, \ell, S)$ is visited by a path, then $x_v \in [\gamma^-_{a,\ell}, \gamma^+_{a,\ell}]$ and $\tau^{\min}_a \leqslant \tau_v \leqslant \tau^{\max}_a$.

- if an arc $e = (v, v')$ labeled with m is visited by a path, then $x_v + \tau_v + d_{\ell,\ell',m} \leqslant x_{v'}$, where ℓ and ℓ' are respectively the locations of v and v'.

- for each agent, the duration of activities and the trips performed should cover the time horizon T.

- for each agent n, the starting time of the first vertex $\texttt{dawn}(n)$ must be equal to 0, and the ending time of the last vertex $\texttt{dusk}(n)$ must be equal to T.

Finally, for each agent n, the cost of activities and trips performed during the day must stay below the budget B_n.

4.2 MILP

To ease the writing of the MILP modeling the problem of Section 4.1, for every vertex v, we denote the triple it represents by (a_v, ℓ_v, S_v). We also define the duration and the cost along an arc $e = (u, v)$ labeled with m_e by $d_e = d_{\ell_u,\ell_v,m_e}$ and $\rho_e = \rho_{\ell_u,\ell_v,m_e}$.

4.2.1 Variables

We introduce two types of decision variables: binary variables, which define the agents' paths through the graph, and continuous variables, which determine the temporal allocation of activities. The binary variables z^n_e indicate whether agent n takes arc e. The binary variables w_v denote whether vertex v is activated, meaning that the path of each agent in S_v goes through v. The continuous variables x_v and τ_v represent, respectively, the starting time and the duration of activity a_v for the agents in S_v at location ℓ_v.

4.2.2 Combinatorial constraints

The following constraints are combinatorial constraints on vertices and arcs.

$$w_v = \sum_{e \in \delta^+(v)} z_e^n \qquad \forall v \in V \quad \forall n \in S_v \tag{4a}$$

$$z_e^n = 0 \qquad \forall e = (u, v) \in E \quad \forall n \notin N_u \cap N_v \tag{4b}$$

$$\sum_{v \in V:\, a_v \in G_k} w_v \geqslant p_k \qquad \forall k \in K \tag{4c}$$

$$w_v + w_{v'} \leqslant 1 \qquad \forall v, v' \in V \text{ s.t. } a_v = a_{v'}, S_v = S_{v'}, \text{ and } \ell_v \neq \ell_{v'} \tag{4d}$$

Constraints (4a) ensure that an activity vertex v is marked as visited (i.e., $w_v = 1$) if and only if each agent in the group S_v selects an outgoing arc from v. This links the binary activity decision variable w_v with the individual arc variables z_e^n. Constraints (4b) prohibit agents from using arcs that are not eligible to use. Constraints (4c) ensure that certain mandatory groups of activities are collectively satisfied. For each group G_k, a minimum number p_k of vertices from the group must be activated. Constraints (4d) ensure exclusivity in location choice: if two vertices v and v' correspond to the same activity and group but are located at different places, at most one of them can be selected.

4.2.3 Time-related constraints

The next constraints are time-related constraints.

$$x_v \geqslant x_u + \tau_u + d_e - M_e(1 - z_e^n) \qquad \forall e = (u, v) \in E \quad \forall n \in N \tag{5a}$$

$$\gamma^-_{a_v,\ell_v} w_v \leqslant x_v \leqslant \gamma^+_{a_v,\ell_v} w_v - \tau_v \qquad \forall v \in V \tag{5b}$$

$$\tau^{\min}_{a_v} w_v \leqslant \tau_v \leqslant \tau^{\max}_{a_v} w_v \qquad \forall v \in V \tag{5c}$$

$$\sum_{v \in V:\, n \in S_v} \tau_v + \sum_{e \in E} d_e z_e^n = T \qquad \forall n \in N \tag{5d}$$

$$x_{\mathrm{dawn}(n)} = 0, \; x_{\mathrm{dusk}(n)} + \tau_{\mathrm{dusk}(n)} = T \qquad \forall n \in N \tag{5e}$$

Constraints (5a) enforce temporal consistency along arcs. If an agent n uses arc $e = (u, v)$, the start of the activity at vertex v cannot occur before the completion of the activity at vertex u and the corresponding travel between them. M_e is the largest constant ensuring the constraint becomes non-binding when the arc is not used. For an arc $e = (u, v)$, its value is $M_e = \gamma^+_{a_u} + d_e$. Constraints (5b) ensure that if an activity is selected at vertex v, its start time x_v falls within the time window $[\gamma^-_{a_v,\ell_v}, \gamma^+_{a_v,\ell_v}]$. If the activity is not selected, the constraint becomes inactive. Constraints (5c) enforce that the duration τ_v of an activity respects its minimum duration $\tau_{a_v}^{\min}$ when the vertex is visited, and becomes inactive otherwise. Constraint (5d) ensure that that the total time spent by any agent performing activities and trips equals the total time period T. Constraint (5e) enforce that each agent starts and ends their schedule with activities `dawn` and `dusk` respectively.

4.2.4 MILP

Now that the combinatorial and temporal constraints have been established, the MILP can be written as follows.

$$
\begin{aligned}
\max_{(z,w,x,\tau)} \quad & U(z,w,x,\tau) && && (6a)\\
\text{s.t.} \quad & \sum_{e\in\delta^+(v)} z_e^n = \sum_{e\in\delta^-(v)} z_e^n && \forall v \in V, \forall n \in N && (6b)\\
& \sum_{e\in\delta^+(\mathtt{dawn}(n))} z_e^n = 1 && \forall n \in N && (6c)\\
& \sum_{e\in\delta^+(\mathtt{dusk}(n))} z_e^n = 1 && \forall n \in N && (6d)\\
& \sum_{v\in V:\, n\in S_v} c_{a_v \ell_v} w_v + \sum_{e\in E} \rho_e z_e^n \leqslant B_n && \forall n \in N && (6e)\\
& \text{Combinatorial constraints} && (4a)\text{–}(4d) &&\\
& \text{Time-related constraints} && (5a)\text{–}(5e) &&\\
& z_e^n \in \{0,1\}, \quad w_v \in \{0,1\} && \forall e \in E, \forall n \in N, \forall v \in V &&
\end{aligned}
$$

where (6a) is the objective function defined later in Section 4.2.5. Constraints (6b) are flow conservation constraints. Constraints (6c) and (6d) ensure that each agent's path starts at their corresponding `dawn` vertex and ends at their `dusk` vertex. Constraints (6e) impose a financial budget B_n on each agent n. It accounts for the cumulative costs of performing activities and traveling along selected arcs.

4.2.5 Objective

The objective function (6a) is a translation of the utility described in Section 3.1. It includes rewards for selected activities, penalties for deviations from preferred start times and durations, penalties for travel cost and travel time, and additional rewards for joint travel between agents. For $e = (v, v')$, $v = (a, \ell, S)$, and $v' = (a', \ell', S')$, we define

$$
r_{en} := r_{\ell\ell' mn} + \sum_{n'\in N} r^{\text{joint}}_{\ell\ell' mn'} z_e^{n'} .
$$

Hence, the objective function is written as follows:

$$
\begin{aligned}
U(z, w, x, \tau) = & \sum_{v \in V} u_v w_v \\
& + \sum_{v=(a,\ell,S) \in V} \sum_{n \in S} \omega_n (\theta_{x\,\text{early}} (x^*_{an} - x_v)^+ + \theta_{x\,\text{late}} (x_v - x^*_{an})^+) \\
& + \sum_{v=(a,\ell,S) \in V} \sum_{n \in S} \omega_n (\theta_{\tau\,\text{short}} (\tau^*_{an} - \tau_v)^+ + \theta_{\tau\,\text{long}} (\tau_v - \tau^*_{an})^+) \\
& + \sum_{e \in E} \sum_{n \in N} \omega_n (u_e + r_{en}) z^n_e
\end{aligned}
$$

In this form, the utility function exactly corresponds to the one defined in Equation (3) of Section 3, and it is straightforward to linearize.

4.3 Back to the example

We now return to the example described in Section 3.3. Assume that a graph G has been constructed following the steps of Section 4.1. First, we show that the vertices and the arcs displayed in Figure 2 form a subgraph H of G. We then provide an interpretation of the feasible paths that pass through this subgraph.

Let H be the graph of Figure 2. For each vertex $v = (a, \ell, S)$, the location ℓ corresponds to a possible location for activity a and the agents of S satisfy the conditions listed in Section 4.1. Moreover, the arcs (v_1, v_3), (v_1, v_4), and (v_3, v_4) are labeled with `car` because it is the only mode available for the agent Car between any couple of different locations. The arc (v_3, v_2) is labeled with `noTrip` because vertices v_2 and v_3 have the same location `city`. Finally, the arc (v_1, v_3) exists and is labeled with `PT`, because Bob can use the public transport to travel from the `office` to `city`. Hence, Figure 2 illustrates a small but relevant subgraph H of the full graph G built as in Section 4.1.

A feasible solution to the multi-agent scheduling problem corresponds to a set of paths in the graph G - one path per agent - and these paths may or may not pass through this specific subgraph H. However, if the paths of Alice, Bob, and Car pass through the subgraph H, there are two possibilities:

- The paths of Alice and her car contain vertices v_1, v_3, and v_4 and the path of Bob contains v_1, v_3, and v_2. It means that Alice escorts Bob to his tennis activity in the city using her car, and then returns home alone.
- The paths of Alice and her car contain vertices v_1 and v_4 and the path of Bob contains

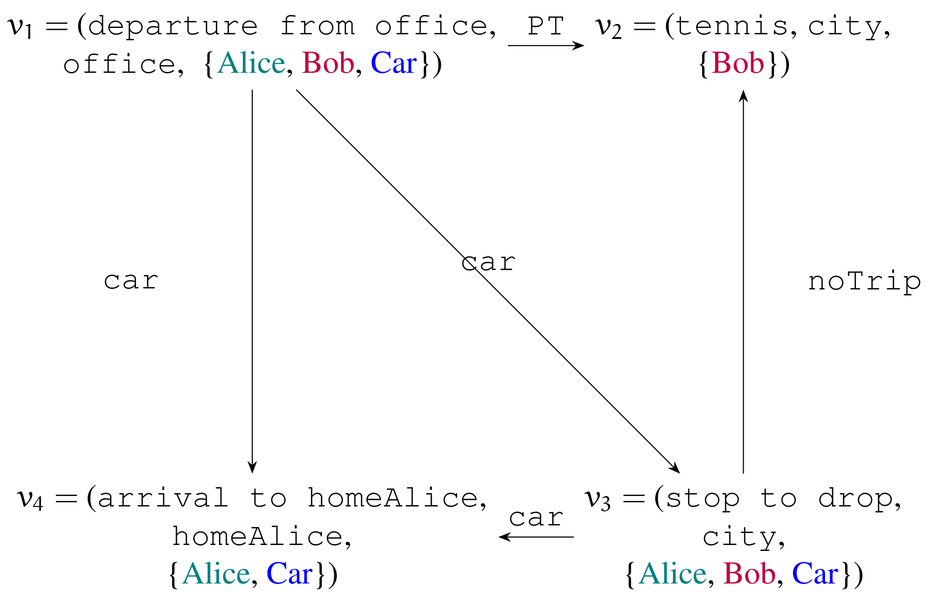

Figure 2: Subgraph illustrating a ride-sharing situation.

v_1 and v_3. The interpretation is that Alice goes home alone and Bob uses the public transport to go to his tennis activity.

Hence, depending on the rewards associated with the activities and the trips, as well as the weights for the deviations from the schedule, and the weights for the travel time in the utility function, Alice and Bob may decide to carpool or not. This example demonstrates that our model effectively captures the trade-offs involved in collective decision-making, specifically in the context of carpooling. This model can easily accommodate other situations as well. For example, escorting someone to an activity, arriving late to a shared activity, or coordinating shared travels on public transport are all scenarios which could enhance the modeling of behavior, while being integrated in an ABM.

5 Numerical experiments

This section presents the results of our numerical experiments, which aim to achieve two objectives. First, we validate our model by comparing its computational performance with the formulation of Rezvany et al. (2023) on real instances. Second, we demonstrate the flexibility of our model by showing that it incorporates new features that are useful to model real situations, that were previously beyond reach.

5.1 Instances and experimental setting

To evaluate both computational performance and flexibility, we constructed two sets of instances. The first set of *restricted instances* ensures full equivalence with the problem solved by Rezvany's formulation. The second set of *extended instances*, built on the very same data, exploits the additional flexibility offered by our formulation to better reflect real-world scenarios (e.g., variable transport modes, and social activities by any subset of agents).

Both MILPs are solved using GUROBI OPTIMIZER v12.0.3. All experiments are conducted on a machine equipped with 72 *Intel(R) Xeon(R) Platinum 8360Y CPU 2.40GHz* processors, running *Red Hat Enterprise Linux 9.4 (Plow)*. The machine has a total of 503GB of available RAM. The time limit for each experiment is set to 600 seconds.

The test instances are derived from the *UK National Travel Survey (NTS)* (Department for Transport, 2024), from which we selected twelve households. The NTS provides information on individuals' activities, locations, preferred schedules, and available transport modes over several days. For our experiments, we aggregate this data to construct a representative set of possible activities within a single day, and we define the time horizon accordingly as one day. For each activity type, consistent opening and closing hours are specified, together with minimum and maximum duration bounds to reflect realistic daily behavior. A detailed description of the data is provided in Table 8 (Appendix).

The parameters used in the utility functions are taken from the literature and their values are summarized in Table 5.

Parameter	**Value**	**Reference**
θ_c	-1 per CHF	Assumed
$\theta_{x\text{ early}}, \theta_{x\text{ late}}$	-2.4 per hour	Pougala et al., 2022
$\theta_{\tau\text{ short}}, \theta_{\tau\text{ long}}$	-2.4 per hour	Pougala et al., 2022
θ_{ttm} for m private	-1 per hour	Pougala et al., 2022
θ_{ttm} for m public	-0.4 per hour	Pougala et al., 2022
θ_{tc}	-1 per CHF	Pougala et al., 2022
$r^{\text{joint}}_{\ell\ell' mnn'}$	0.5	Assumed
$r^{\text{joint}}_{ann'}$	1	Rezvany et al., 2025
$r_{a\ell n}$	From 5 to 15 depending on the activity	Rezvany et al., 2025

Table 5: Parameters used in the experiments.
"CHF" refers to the Swiss Franc.

5.2 Results

5.2.1 Model validation: comparison with the previous formulation

The left-hand side of table 6 compares the performance of the proposed model with the reference implementation by Rezvany et al. (2023) across the twelve instances, called restricted instances, described in the previous subsection. For each instance, the computational time, optimality gap, and problem size in terms of number of agents ($|N|$) and activities ($|A|$) are reported. Our approach reaches optimality on all instances, while the reference model fails to do so on the larger ones (instances 9–10). In addition to guaranteeing optimality, the proposed formulation is substantially faster: on average, the model of Rezvany et al. (2023) is about 9.7

times slower than ours on this set of instances. Overall, the graph-based approach substantially improves computational efficiency.

	|N|	|A|	Restricted inst.		Rezvany et al.			Extended inst.	
			Time	**Gap**	**Time**	**Gap**	**Ratio**	**Time**	**Gap**
H 1	2	20	4.76	0%	2.14	0%	0.45	6.54	0%
H 2	4	42	10.54	0%	276.03	0%	26.2	600.00	12.53%
H 3	2	32	4.21	0%	11.20	0%	2.66	18.26	0%
H 4	1	10	0.32	0%	1.14	0%	3.56	0.34	0%
H 5	1	16	0.62	0%	4.97	0%	8.02	0.62	0%
H 6	2	16	0.64	0%	1.14	0%	1.78	0.89	0%
H 7	3	26	9.01	0%	70.20	0%	7.79	534.32	0%
H 8	2	34	6.07	0%	43.21	0%	7.12	41.55	0%
H 9	4	45	364.91	0%	600.00	133.7%	1.65	600.00	2.16%
H 10	2	42	496.81	0%	600.00	66.2%	1.21	600.00	0.58%
H 11	2	32	4.57	0%	246.64	0%	54.0	42.45	0%
H 12	2	11	0.64	0%	0.93	0%	1.45	0.78	0%

Table 6: Comparison of solving time and optimality gap between our model and Rezvany et al. (2023) on the restricted instances, including the extended instances.

Table 6 also compares the performance of our model on both restricted and extended instances. For the same households, we observe that the inclusion of more features results in a significant increase in computational time. Mathematically, solving the extended instance is akin to solving a relaxation of the corresponding restricted instance. However, this increased complexity allows our model to capture more detailed behaviors, as explained in the next subsection.

5.2.2 Model flexibility: capturing new behaviors

To demonstrate the flexibility of our model, we use the set of extended instances built from the same households, but allowing for variable transport modes and social activities performed by any subset of agents. Table 7 compares the optimal solutions obtained on the restricted instances, corresponding to Rezvany's formulation, and on the extended instances. Differences in the results are observed in 5 out of 12 cases, confirming that the two formulations can lead to distinct optimal schedules. The additional behaviors captured by our model are listed in the column "Differences." They include: (different modes) agents using different transport modes after a shared activity, (staggered social activity) agents participating in shared activities with staggered timing—some arriving later to accommodate individual activities, and (not all agents for social activities) agents not performing shared activities without preventing others from performing them. These patterns illustrate that our formulation can capture a broader and more realistic range of coordination behaviors within households.

	$\|N\|$	$\|A\|$	**Differences**
Household 1	2	20	(different modes)
Household 2	4	42	(staggered social activity)
Household 3	2	32	(staggered social activity), (not all agents for social activities)
Household 4	1	10	-
Household 5	1	16	-
Household 6	2	16	-
Household 7	3	26	(staggered social activity)
Household 8	2	34	-
Household 9	4	45	-
Household 10	2	42	-
Household 11	2	32	-
Household 12	2	11	(not all agents for social activities)

Table 7: Comparison of the restricted and extended instances.

To further illustrate the behavioral implications of the additional flexibility, we conduct a simulation experiment on a representative instance (Household 7). This household includes three agents: Agent 1 performs an activity *Home*, Agent 2 engages in activities *work* and *shopping*, and Agent 3 attends an activity *education*. We can assume that Agents 1 and 2 are parents and that Agent 3 is a schoolchild. The social activities consist of *Escort education* (which, on the extended instances, can be performed by any subset of two agents including the schoolchild), *Other social* (a social activity outside home that can involve any subset of agents), and *Visit friends/relatives at home* (a social activity at home, also feasible for any subset of agents). In contrast, in the restricted model, all these social activities must either be performed jointly by all household members or not at all.

We draw 100 different sets of error terms $(\xi_{an})_{n\in N,\, a\in A_n}$ and $(\xi_{\ell\ell' mn})_{n\in N,(\ell,\ell')\in L^2, m\in M_n^{\ell,\ell'}}$, each sampled independently from a standard normal distribution. For each draw, we solve both the restricted and extended instance to obtain the optimal schedules.

Figure 3 reports the distribution of activities over the course of the day across these simulated optimal schedules, using the same error terms for both models. The results highlight substantial behavioral differences between the two formulations. At any given time of day, the extended instances produce a wider variety of feasible activity configurations across agents, indicating greater variability in the simulated optimal schedules. Moreover, on the extended instances, activities such as *Escort education* can be performed by different subsets of agents including the schoolchild, while social activities like *Other social* and *Visit friends/relatives at home* may involve only part of the household or asynchronous participation—where one agent joins or leaves the shared activity at a different time. These patterns are impossible in the restricted formulation, where all shared activities must start and end simultaneously for all participants.

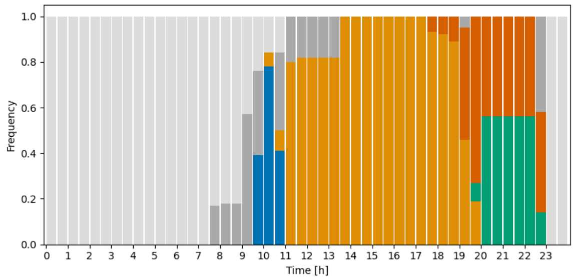

(a) Agent 1 – Extended

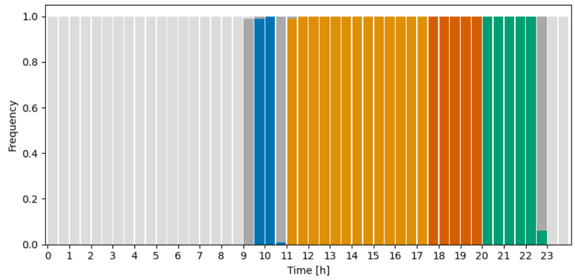

(b) Agent 1 – Restricted

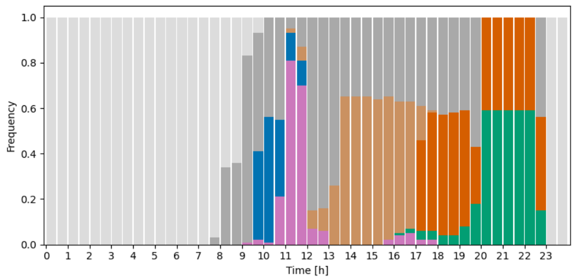

(c) Agent 2 – Extended

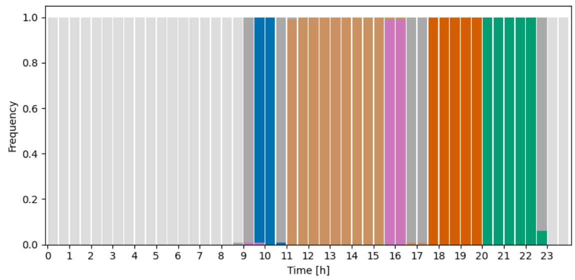

(d) Agent 2 – Restricted

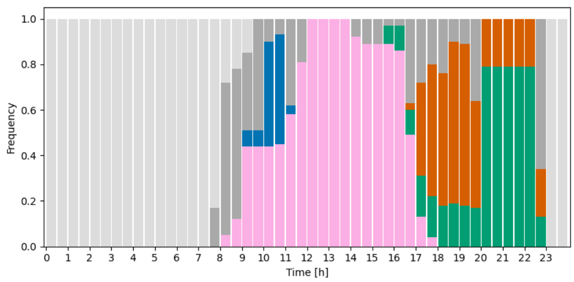

(e) Agent 3 – Extended

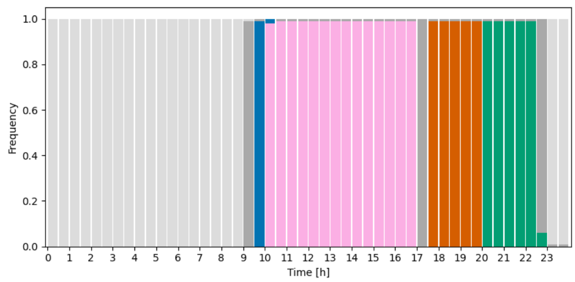

(f) Agent 3 – Restricted

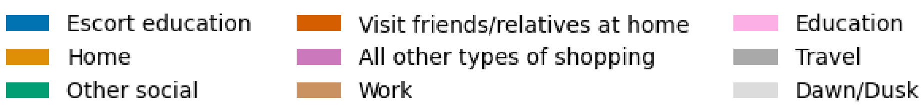

Figure 3: Distribution of activities during the day for Household 7's members, comparing the restricted (left) and extended (right) instances.

Figure 4 presents the optimal schedules for one draw from the simulation. These schedules illustrate the differences in agent activity timings between the two formulations. In the restricted instances (panels a, b, and c), social activities are strictly synchronized, with all participants starting and ending at the same time. In contrast, the extended instances (panels d, e, and f) allow for more flexibility, where agents can perform activities asynchronously, with some activities being performed by a subset of agents.

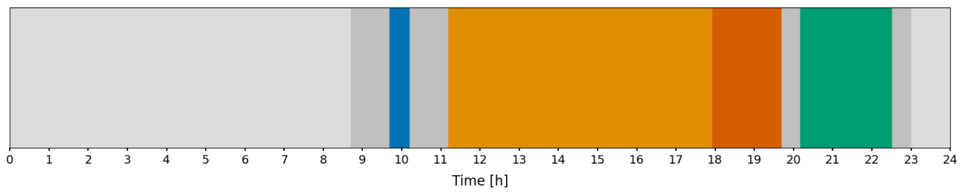

(a) Agent 1 – Restricted

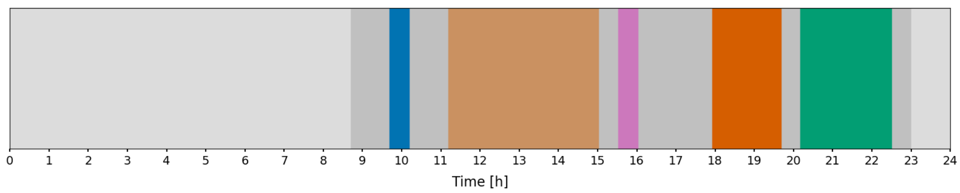

(b) Agent 2 – Restricted

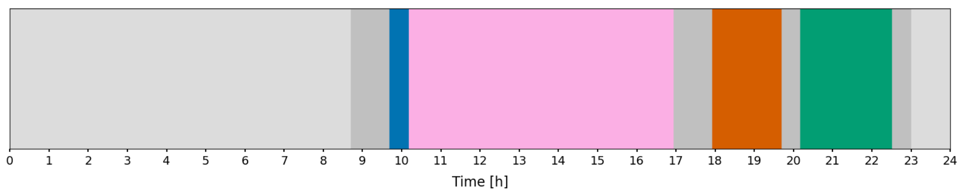

(c) Agent 3 – Restricted

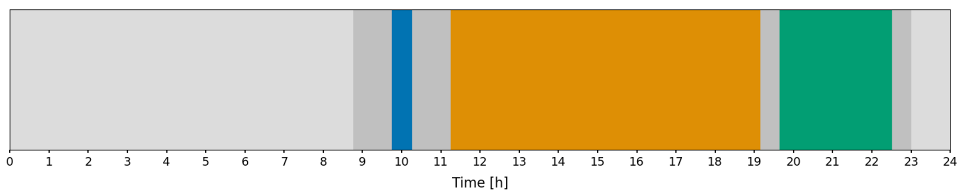

(d) Agent 1 – Extended

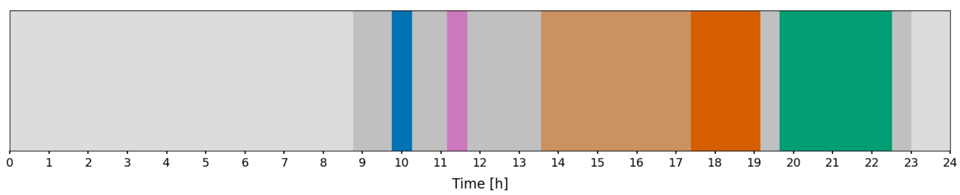

(e) Agent 2 – Extended

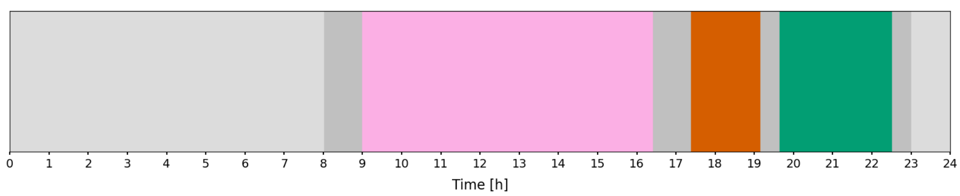

(f) Agent 3 – Extended

Figure 4: Optimal schedule of one draw from the simulation for the restricted (above) and extended (below) instances.

5.3 Discussion

The results demonstrate that our model not only outperforms the formulation of Rezvany et al. (2023) in terms of computational efficiency but also offers greater flexibility to capture real-world behaviors. The improved performance can be explained by several structural differences. First, shared activities are represented directly within the event–activity graph through vertices that specify the subset of participating agents. Second, the flow constraints are expressed compactly, and only a single big-M constraint is required. Third, transport modes are modeled as attributes of arcs—rather than by duplicating each activity once per possible departure mode. Together, these choices lead to faster convergence on most instances.

Despite this improvement, Table 7 shows that the computational time increases when the full flexibility of our model is exploited. This is expected, since our model constitutes a relaxation of the formulation of Rezvany et al. (2023). The complexity mainly arises from joint activities: in their absence, the problem can be decomposed into independent single-agent scheduling problems that can be solved in parallel. These observations suggest that the next step should focus on algorithmic strategies that exploit the graph structure—such as dynamic programming, or Lagrangian relaxation—to handle the coupling introduced by joint activities more efficiently.

Beyond computational aspects, the results also highlight the behavioral richness enabled by the proposed formulation. The fact that optimal schedules differ in several real instances confirms that this flexibility is not merely theoretical but also valuable in practice. In the experiments, we already introduced flexibility by allowing different transport modes after a shared activity and by enabling different subsets of agents to participate in shared activities. Yet, the same formulation could easily accommodate additional features—such as groups of mandatory activities or participation limits—further enhancing the realism of the resulting schedules.

Finally, the same framework can be extended beyond households to other coordinated groups. The concept of non-human resources can be generalized to represent shared elements such as electric vehicle chargers, meeting rooms, or workplaces, enabling the model to capture synchronization among co-workers or other small collectives. The main limitation for applying these multi-agent models to broader contexts remains data availability, rather than modeling capability.

6 Concluding remarks

This paper addresses the challenge of representing coordination and synchronization among multiple agents within ABMs. We proposed a new graph-based mixed-integer linear formulation that generalizes previous approaches and enables the modeling of a wider range of realistic behaviors. Built upon a labeled graph, the formulation produces interpretable solutions—one path per agent in the graph—while reducing redundancy and preserving tractability.

The numerical experiments on UK National Travel Survey households demonstrated that the model consistently improves computational efficiency and realism compared to existing formulations. More importantly, the results confirmed that flexibility in agent participation and transport modes meaningfully affects the resulting optimal schedules, highlighting the importance of capturing such interactions in multi-agent settings.

Beyond its direct results, this work establishes a foundation for a new generation of optimization-driven ABMs. The formulation can naturally incorporate additional behavioral features—for example, modeling the shared use of new resources or limiting the number of times certain activities may be performed by the group, and it opens the door to new algorithmic developments that exploit its structure. Future research will focus on decomposition and relaxation techniques to improve scalability, and on extending the framework to broader coordination contexts, such as co-working or shared mobility systems.

References

Adler, T. and Ben-Akiva, M. (1979). A theoretical and empirical model of trip chaining behavior, *Transportation Research Part B: Methodological* **13**(3): 243–257.

Axhausen, K. W. and Gärling, T. (1992). Activity-based approaches to travel analysis: conceptual frameworks, models, and research problems, *Transport Reviews* **12**(4): 323–341.

Bhat, C. R. (2005). A multiple discrete–continuous extreme value model: formulation and application to discretionary time-use decisions, *Transportation Research Part B: Methodological* **39**(8): 679–707.

Bredström, D. and Rönnqvist, M. (2008). Combined vehicle routing and scheduling with temporal precedence and synchronization constraints, *European Journal of Operational Research* **191**(1): 19–31.

Bundesamt für Statistik (BFS) (2025). Kosten und finanzierung des verkehrs 2022, *Technical report*, Federal Statistical Office, Neuchâtel, Switzerland.

Chapin, F. S. (1974). *Human Activity Patterns in the City: Things People Do in Time and in Space*, Wiley-Interscience.

Cirillo, C. and Axhausen, K. W. (2006). Evidence on the distribution of values of travel time savings from a six-week diary, *Transportation Research Part A: Policy and Practice* **40**(5): 444–457.

De Palma, A., Picard, N. and Inoa, I. (2014). Discrete Choice Decision-Making with Multiple Decision Makers within the Household, *hal-00969216* .

Department for Transport (2024). National travel survey: 2002–2023. SN: 7553, distributed by the UK Data Service.

Drexl, M. (2012). Synchronization in Vehicle Routing—A Survey of VRPs with Multiple Synchronization Constraints, *Transportation Science* **46**(3): 297–316.

Ettema, D., Bastin, F., Polak, J. and Ashiru, O. (2007). Modelling the joint choice of activity timing and duration, *Transportation Research Part A: Policy and Practice* **41**(9): 827–841.

Gan, L. P. and Recker, W. (2013). Stochastic Preplanned Household Activity Pattern Problem with Uncertain Activity Participation (SHAPP), *Transportation Science* **47**(3): 439–454. Publisher: INFORMS.

Gaul, D., Klamroth, K. and Stiglmayr, M. (2022). Event-based MILP models for ridepooling applications, *European Journal of Operational Research* **301**(3): 1048–1063.

Greeven, P., Jara-Díaz, S. R., Munizaga, M. A. and Axhausen, K. W. (2005). Calibration of a joint time assignment and mode choice model system, p. 27.

Guarda, P. and Qian, S. (2024). Statistical inference of travelers' route choice preferences with system-level data, *Transportation Research Part B: Methodological* **179**: 102853.

Gupta, S., Vovsha, P. S., Livshits, V., Maneva, P. and Jeon, K. (2014). Incorporation of escorting children to school in modeling individual daily activity patterns of household members, *Transportation Research Record* **2429**: 20–29.

Habib, K. M. N. and Hui, V. (2017). An activity-based approach of investigating travel behaviour of older people, *Transportation* **44**(3): 555–573.

Kurata, L. and Nakamura, K. (2025). Collective decisions under uncertainty: efficiency, ex-ante fairness, and normalization, *arXiv preprint arXiv:2505.03232* .

Lai, X., Lam, W. H. K., Su, J. and Fu, H. (2019). Modelling intra-household interactions in time-use and activity patterns of retired and dual-earner couples, *Transportation Research Part A: Policy and Practice* **126**: 172–194.

Manser, P., Haering, T., Hillel, T., Pougala, J., Krueger, R. and Bierlaire, M. (2024). Estimating flexibility preferences to resolve temporal scheduling conflicts in activity-based modelling, *Transportation* **51**(2): 501–528.

Palma, D., Enam, A., Hess, S., Calastri, C. and Crastes dit Sourd, R. (2021). Modelling multiple occurrences of activities during a day: an extension of the MDCEV model, *Transportmetrica B: Transport Dynamics* **9**(1): 456–478. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/21680566.2021.1900755.

Pawlak, J., Faghih Imani, A. and Sivakumar, A. (2021). How do household activities drive electricity demand? Applying activity-based modelling in the context of the United Kingdom, *Energy Research & Social Science* **82**: 102318.

Pougala, J., Hillel, T. and Bierlaire, M. (2022). Capturing trade-offs between daily scheduling choices, *Journal of Choice Modelling* **43**: 100354.

Pougala, J., Hillel, T. and Bierlaire, M. (2023). OASIS: Optimisation-based Activity Scheduling with Integrated Simultaneous choice dimensions, *Transportation Research Part C: Emerging Technologies* **155**: 104291.

Recker, W. W. (1995). The household activity pattern problem: General formulation and solution, *Transportation Research Part B: Methodological* **29**(1): 61–77.

Rezvany, N., Bierlaire, M. and Hillel, T. (2023). Simulating intra-household interactions for in- and out-of-home activity scheduling, *Transportation Research Part C: Emerging Technologies* **157**: 104362.

Rezvany, N., Hillel, T. and Bierlaire, M. (2025). Household activity scheduling: choice-set generation and parameter estimation, *TRANSP-OR Report 250712*, Transport and Mobility Laboratory, École Polytechnique Fédérale de Lausanne.

Roorda, M. J., Miller, E. J. and Habib, K. M. N. (2008). Validation of TASHA: A 24-h activity scheduling microsimulation model, *Transportation Research Part A: Policy and Practice* **42**(2): 360–375.

Shilov, I., Elokda, E., Hall, S., Nax, H. H. and Bolognani, S. (2025). Welfare and cost aggregation for multi-agent control: When to choose which social cost function, and why?, *arXiv preprint arXiv:2503.20772* .

Srinivasan, K. K. and Athuru, S. R. (2005). Analysis of within-household effects and between-household differences in maintenance activity allocation, *Transportation* **32**(5): 495–521.

Vo, K. D., Lam, W. H. K., Chen, A. and Shao, H. (2020). A household optimum utility approach for modeling joint activity-travel choices in congested road networks, *Transportation Research Part B: Methodological* **134**: 93–125.

Vovsha, P., Bradley, M. and Bowman, J. L. (2005). Activity-based travel forecasting models in the united states: Progress since 1995 and prospects for the future, *Progress in Activity-Based Analysis*, Elsevier, Kidlington, Oxford, United Kingdom, pp. 389–414.

Vovsha, P., Petersen, E. and Donnelly, R. (2003). Explicit Modeling of Joint Travel by Household Members: Statistical Evidence and Applied Approach, *Transportation Research Record* **1831**(1): 1–10. Publisher: SAGE Publications Inc.

Xu, Z., Kang, J. E. and Chen, R. (2018). A random utility based estimation framework for the household activity pattern problem, *Transportation Research Part A: Policy and Practice* **114**: 321–337.

Appendix A

Requirement	Acquired solution
Set of agents N	Members of a household of the dataset
Set of activities A_n	Generated from the activities of the actual schedule of agent n over several days. The social activities are extended to all the household
Set of locations L_a	Actual location of activity a in the dataset
$c_{a,\ell}$	Set to 0 for the experiments
G_k	Activities of the same type can only be done once per agent
C_a	Set to $\|N\|$ to be non restrictive
N_a	For a social activity a in the restricted instance, $N_a = \{N\}$, and $N_a = \varnothing$ otherwise
Set of transport modes M_n^{kl}	For $k \neq \ell$, aggregation of the modes used by n in the dataset, for $k = \ell$, equals to $\{\texttt{walk}\}$
Travel cost $\rho(k, \ell, m)$	Proportional to the travel distance (geographic distance computed from the actual locations), with coefficient taken from (Bundesamt für Statistik (BFS), 2025)
Travel time $d(k, \ell, m)$	Proportional to the travel distance, with a coefficient taken from (Bundesamt für Statistik (BFS), 2025)
Preferred start time $x^{\star}_{an}$ and duration $\tau^{\star}_{an}$	Based on recorded values in the dataset
Feasible time window $[\gamma^-_{a\ell}, \gamma^+_{a\ell}]$	Set depending on the type of activity (e.g., shopping from 9:00 to 20:00)
Minimum and maximum activity duration $[\tau^{\min}_a, \tau^{\max}_a]$	Minimum 30 minutes, and maximum 24 hours (non constraining)
Time budget T	Set to 24 hours
Financial budget B_n	Set to a high value to be non restrictive
Agent priority ω_n	Set to 1 for all agents

Table 8: Model data requirements used in the experiments.

Appendix B

Algorithm 1 Vehicle-based Auxiliary Activity Generation Algorithm

Input: Set N of agents, drivers $D \subseteq N$, vehicle $v \in N$, activity sets $\{A_n \colon n \in N \setminus \{v\}\}$, empty set A_v, location sets $\{L_a \colon a \in A_n\}$

Output: Updated activity sets $\{A_n \colon n \in N\}$, location sets $\{L_a \colon a \in A_v\}$, and required-agent sets $\{N_a \colon a \in A_v\}$ for the activities of vehicle v

for each agent $n \in N \setminus \{v\}$ **do**
    for each activity $a \in A_n$ **do**
        Create activities $\texttt{arrival\_to}(a)$ and $\texttt{departure\_from}(a)$
        Set $L_{\texttt{arrival\_to}(a)} \leftarrow L_a$ and $L_{\texttt{departure\_from}(a)} \leftarrow L_a$
        Add $\texttt{arrival\_to}(a)$ and $\texttt{departure\_from}(a)$ to:

$$A_n,\ A_v,\ \text{and } A_d \text{ for each } d \in D.$$

        Set required-agent sets:

$$N_{\texttt{arrival\_to}(a)} \leftarrow \big\{ \{v, d\} \colon\ d \in D \big\}, \qquad N_{\texttt{departure\_from}(a)} \leftarrow \big\{ \{v, d\} \colon\ d \in D \big\}.$$

        if $n \in D$ **then**
            Create activity $\texttt{parking}(a)$ with $L_{\texttt{parking}(a)} \leftarrow L_a$
            Add $\texttt{parking}(a)$ to A_v
        end if
    end for
end for
return $\{A_n\}$ for all $n \in N$, and $\{L_a, N_a\}$ for all $a \in A_v$

For all activities created by the algorithm (`arrival_to_a`, `departure_from_a`, and `parking_a`), the following parameters must also be initialized to ensure full model consistency:

$$c_{a,\ell} = 0, \qquad C_a = 4, \qquad x^{\star}_{an}, \quad \tau^{\star}_{an}, \quad [\gamma^{-}_{a\ell}, \gamma^{+}_{a\ell}], \quad [\tau^{\min}_a, \tau^{\max}_a].$$

The parameters $x^{\star}_{an}$, $\tau^{\star}_{an}$, $[\gamma^{-}_{a\ell}, \gamma^{+}_{a\ell}]$, and $[\tau^{\min}_a, \tau^{\max}_a]$ are defined consistently with the corresponding activity a from which the auxiliary activities are derived. Their specific initialization is not constrained by the algorithm itself.