Designing Layouts for Sequential Experiences: Application to Cultural Institutions

Ali Aouad[1] Abhishek Deshmane[2] Victor Martínez-de-Albéniz[3]

Abstract

A fundamental issue faced by experience providers – ranging from retailing to cultural institutions –is to display a collection of items for physical or digital interactions. The arrangement of the exhibits in different locations, which we call the *layout*, affects the visitors' choices over time and space, thereby driving their engagement with the offered experience. This paper develops a data-driven analytics framework to inform such operational decisions, taking into account visitors' preferences. First, we propose a dynamic choice model, called *Pathway MNL*, that represents visitor activity as a sequence of conditional logit experiments influenced by the layout. We estimate this model on large-scale data logs of multimedia guide usage at the Van Gogh Museum (Netherlands). Using parametric specifications of the utility function, we uncover significant relationships between visitors' choices and layout distances, artwork characteristics, and other contextual dimensions. Visitors value proximity and variety when locally constructing their path into the museum, but their choices are also influenced by the level of congestion and the number of artworks already seen. Our model predicts the next visitor action with an out-of-sample accuracy of 63%. Natural experiments on the layout provide further empirical validations. Second, we formulate the *layout optimization* problem, where the goal is to assign artworks to different locations to maximize the expected length of visitors' paths. We establish a strong inapproximability result for this new optimization setting. We identify realistic interventions that can significantly lift visitors' engagement by improving the attractiveness and retention exercised by the layout.

Keywords: Consumer paths, Sequential choice, Assignment problems, Museums, Cultural operations.

Submitted: July 7, 2022

[1]aaoaud@london.edu. London Business School, Regent's Park, London, United Kingdom.
[2]adeshmane@iese.edu. IESE Business School, University of Navarra, Av. Pearson 21, 08034 Barcelona, Spain.
[3]valbeniz@iese.edu. IESE Business School, University of Navarra, Av. Pearson 21, 08034 Barcelona, Spain.

1 Introduction

Customer experiences take place in physical and virtual spaces that respond to functional needs and offer sensorial, emotional, or intellectual interactions. Individuals navigate these spaces and receive stimuli by engaging with the offerings, which in turn generate utility or disutility, i.e., satisfaction or costs in the form of effort, attention or discomfort. These experiences are central to many economically and culturally fundamental sectors that reach billions of individuals every year, such as retailing, e.g., stores are an essential component to engage with shoppers in an omnichannel world (Gallino and Moreno 2019), entertainment, e.g., theme parks are engineered to bring the thrill-seekers to a "Magic Kingdom" (West 2011), and cultural institutions, e.g., museums seek to leave a lasting memory and an educational footprint on their visitors. In recent years, such physical interactions are often augmented with digital experiences. For example, in retailing, the digital trace of customer trajectories across stores enables targeted promotions based on physical movement (Ghose et al. 2019). Cultural institutions increasingly rely on multimedia content, ranging from audio tracks to augmented reality applications, to construct dynamic and interactive narratives for their art collections (Giannini and Bowen 2019). In other words, experiences are increasingly offered in a hybrid layout, combining the physical space with a digital offering.

The design of the interface between the audience and the offered experiences is a complex task. One must consider that the physical architecture has strong effects on individual behaviors, and more specifically on how agents choose their idiosyncratic paths over time and space. In this sense, experience designers should think like architects that design the space to generate the best possible experience for their users. Even if in some contexts they can fully engineer the experience (see Das Gupta et al. 2016 for instance), they should generally avoid imposing inconvenient paths: most of the time, these are replaced by desire paths, which provide shortcuts from the planned paths designed by the architect. Cases in point are Ikea stores that are seemingly built in maze-like structures with the intention of making the customers go through the entire store assortment. However, shortcuts, which often remain hidden from untrained eyes, are also integrated into these layouts to allow the shoppers in a hurry to skip some parts of the showroom. It is thus not realistic to strictly impose a given trajectory, yet it is worthwhile to understand the connection between the structure of the layout and the realized paths so that the layout can be engineered to improve the experience outcomes. The layout of a digital platform similarly influences what content the users consider and choose to engage with. For example, numerous studies based on clickstream data have established the importance of display positions

Figure 1: Visitor behavior at the Louvre. Source: The authors.

in e-commerce settings (Montgomery et al. 2004, Ursu 2018).

In the cultural sector, the visitor experience in a museum is a rich case study for layout design. The encounter with an artwork is influenced by the notoriety of the artist who created it, its subject matter, its significance in art history, etc. Beyond the aesthetic experience, visitor engagement is shaped by digital and physical layouts. Museums increasingly deploy multimedia guide apps to provide more context about their collection and to recommend certain exhibits (Lanir et al. 2013). Moreover, the placement of the artworks in the rooms and the hanging arrangement influence the exploration of the museum collection. For example, many studies have shown that accessible artworks, located near entrances and exits, attract substantially more traffic (Nielsen 1946); visitors spend more time in front of the first few exhibits they view (Melton 1972, Serrell and Aquarium 1977). Displaying two artworks nearby encourages visitors to view them successively, which facilitates drawing parallels or contrasts between the artworks. To illustrate the dramatic impact layout can have on congestion and visitor attention, Figure 1 shows the *Salle des États* at the Louvre in Paris: on one side of the room, there is a long queue to view Leonardo da Vinci's *Mona Lisa*, while on its opposite side, Veronese's *The Wedding Feast at Cana*, a massive 6.8 m $\times$ 10 m painting that inspired Jacques-Louis David, Eugène Delacroix, and Vincent van Gogh (Greenberger 2021), hardly gets any attention. While the layout has a holistic impact on the visitor journey, most cultural institutions currently focus on curatorial and pedagogical objectives. In museums, artworks are generally displayed by theme, chronology, art movement, etc., to facilitate the contextualisation of the collection during the visitor journey. Nevertheless, visitors' paths often deviate from the narrative constructed by the museum: "Despite great efforts on the part of design teams, it is well documented that many visitors do not view the exhibits in the intended order" (Falk 2016).

The unprecedented availability of data is transforming how the above-mentioned experience providers operate. Retailers, entertainment providers, and cultural institutions now collect data using various technologies, including RFID, WiFi, and mobile devices. This data allows us to track customers in different steps of their journey (Hui et al. 2009b). Yet, there has been little

progress in (1) modeling the relationship between the layout and the visitor pathways, and (2) helping managers design better layouts to serve the curators' experiential objectives. Can such data be leveraged to construct quantitative models for sequential experiences? What layout factors affect visitor pathways? How should a collection of items be assigned to a network of locations, which differ in terms of access and prominence, to maximize visitor engagement?

In this paper, we propose a data-driven framework to answer these questions and support the design of layouts. Our methodological approach is inspired by choice-based revenue management. The trajectory of a visitor is modelled as a pathway of sequential choices. This approach leads to a simple form of dynamic discrete choice model. Consequently, the arrangement of items is formulated as a display optimization problem across a physical network. Through a collaboration with a major museum, we gain insights into the structure of visitor pathways and simulate optimized layout interventions. Our contributions are further detailed below.

Discrete choice modeling for pathways. Our starting point is to view the visitors' movement as a sequential utility-based choice process. We assume that the visitor path is constructed by comparing the underlying utilities of the different options offered to the visitors in each transition (e.g., moving to the next item, skipping, changing floors). Under this assumption, the estimation of the resulting choice model from observed paths reveals the structure of the utility function, which depends on item characteristics and location specifics, maximized by the visitors. For this purpose, we develop a Markov chain-based model, called *Pathway MNL*, where each transition between artworks is captured by a Multinomial Logit (MNL) choice model. In the context of museums, we provide a discussion on the hypothesized effects of artwork, layout, and contextual factors from behavioral theories (Stokols 1976, Lu et al. 2013).

Empirical validations and insights. We collaborate with the Van Gogh Museum in Amsterdam, which is an emblematic cultural institution and one of the most visited museums in the Netherlands. This museum hosts the most comprehensive collection of paintings by Vincent van Gogh (1853-1890), counting over 4,600 artworks and relics. We have gained access to large-scale datasets containing detailed logs of the multimedia tour (MMT) guide usage by 1.5 million visitors for the years 2019-2021. In complement, we constitute a record of the evolution of the physical layout during this period and extract publicly available information about each artwork in the museum collection. By combining these data sources, we fit an econometric specification of the Pathway MNL model and study the relationships between various factors and visitor pathways. As one might expect, the physical distances between two artworks, including walking distances and changes in rooms and floors, strongly reduce the propensity for visitors

to transition from one to another. Similarly, we find that highlighted recommendations and display positions on the multimedia guide significantly affect the visitors' choices. Additionally, we uncover a time pressure effect: as the number of artworks already seen by visitors increases, they are more likely to transition towards the masterpieces of the museum collection. Finally, we find that as the congestion within the museum increases, visitors are more likely to direct themselves towards additional artworks, leading to an increase in the total length of their visit. This suggests that congestion begets congestion, as visitors might not only spend a long time viewing the artworks due to the possible queue but also see more artworks in a congested environment. These findings are followed by a series of out-of-sample prediction comparisons. By leveraging natural experiments during which the layout was modified, we conduct a bias analysis of our model, where we evaluate its ability to predict the effect of new layout changes.

Layout optimization. Using the fitted Pathway MNL model, we simulate counterfactual interventions whereby we modify the assignment of artworks to layout locations. To construct optimized museum layouts, we formulate the artwork-to-location assignment problem as a mathematical program, dubbed the *layout optimization* problem. The objective is to maximize *visitor engagement*, measured through the expected number of artworks clicked on by visitors along their path. While formally showing that the resulting optimization setting is APX-hard, we implement and compare several intuitive policies and a simple heuristic. We find that placing artworks with high attractiveness (according to their 'clickability', a notion similar to that in Besbes et al. 2016 which we formalize later on) in locations far away from the exit can deliver an engagement level better than that of the current layout. Moreover, a few local exchanges of artwork locations can significantly enhance the overall engagement against the current baseline. Interestingly, these layout interventions not only increase the level of interaction with the most accessible artworks, but also generate positive spillovers on nearby options.

The rest of the paper is arranged as follows. We discuss the related literature and our positioning within it in Section 2. This is followed by Section 3 where we explain the context and data. Section 4 has our theoretical model and hypotheses while Section 5 elaborates on the estimation procedure, results, and validation tests. Finally, we provide counterfactuals for enhanced layouts in Section 6 and conclude in Section 7.

2 Related Literature

Pathway data and operations management. Our modeling approach is directly connected to the study of customers' behavior in sequential experiences (Baucells and Sarin 2010, Das Gupta et al. 2016, Bernstein and Martínez-de Albéniz 2017) as well as the incorporation of location effects into operational decision-making (Ahmadi 1997, Rajaram and Ahmadi 2003, Abeliuk et al. 2016, Gallego et al. 2020, Aouad and Segev 2020, Guo et al. 2021). Research on path analysis, with works such as the paper by Thiesse and Fleisch (2008), has relied on the usage of RFID (radio frequency identification) technology such as the PathTracker® system described by Sorensen (2003), with applications to in-store retailing (Hui et al. 2009a,b, 2013). In the context of web browsing and e-commerce, clickstream data is the counterpart source of path information (Montgomery et al. 2004, Besbes et al. 2016). In contrast, the MMT guides offer a hybrid form of data, which tracks movement in a physical layout, but uses clicks and other session-level events to determine the progress of the visitor in the museum. We aim to leverage the granularity offered by our data to understand the visitor preferences and artwork engagement patterns within the museum. Nevertheless, our modeling approach is likely to apply to other sources of information such as RFID and WiFi tracking.

Visitor studies. There has been considerable research in the social sciences on visitor studies to better understand visitor behavior and to inform new curatorial practices. The "Exhibit, Visitor, and Setting" perspectives are proposed by Falk et al. (1985) to inform the design of museum layouts. In particular, the interplay between the physical arrangement, i.e., how artworks are displayed across the exhibition space, and visitors' preferences shape their journey in the museum space (Bourdeau and Chebat 2001). Moreover, visitors can be segmented into different categories based on the way they interact with the artworks, see, e.g., Véron and Levasseur (1989) or Brown and Ratzkin (2011) for the performing arts. Despite numerous behavioral and empirical findings, the literature so far provides very few comprehensive quantitative models to predict and optimize visitor pathways (Güler 2009, Tröndle et al. 2014).

The application of operations research methods to cultural institutions is a largely uncharted area of research. As a rare exception, Martínez-de Albéniz and Valdivia (2019) demonstrates the value of adopting an optimization model to schedule and synchronize the yearly programme of exhibitions. By contrast, our work focuses on the layout and display allocation of the artworks across the museum to help curators and designers improve the visitor experience. There is also a recent line of work that leverages computational methods in traffic management and pedestrian dynamics to predict visitor traffic in museums (Pluchino et al. 2013, Centorrino et al. 2021).

The approach developed in this paper differs by providing a utility-based representation of the visitor choice-making process.

3 Context and Data

3.1 Museum layout

The permanent collection of the Van Gogh Museum, which is the focus of our study, is housed in a building with four levels including the ground floor. The visitors enter the museum through the main door that leads to the ground floor while the subsequent floors can be accessed both by staircases and lifts. Each floor is divided into smaller rooms by the means of walls and partitions which help distinguish different themes. Specifically, the artworks are arranged according to 14 such themes (for example, "Face to Face with Van Gogh", "Painter of Peasant Life", "New Perspectives", etc.) which are distributed across the four levels. Each theme comprises a collection of artworks related to a specific subject matter or period of Van Gogh's life, which altogether convey a dynamic narrative about the evolution of Van Gogh's work. Figure 2a shows the physical layout across the floors, while Figure 2b shows the themes displayed on floor 1.

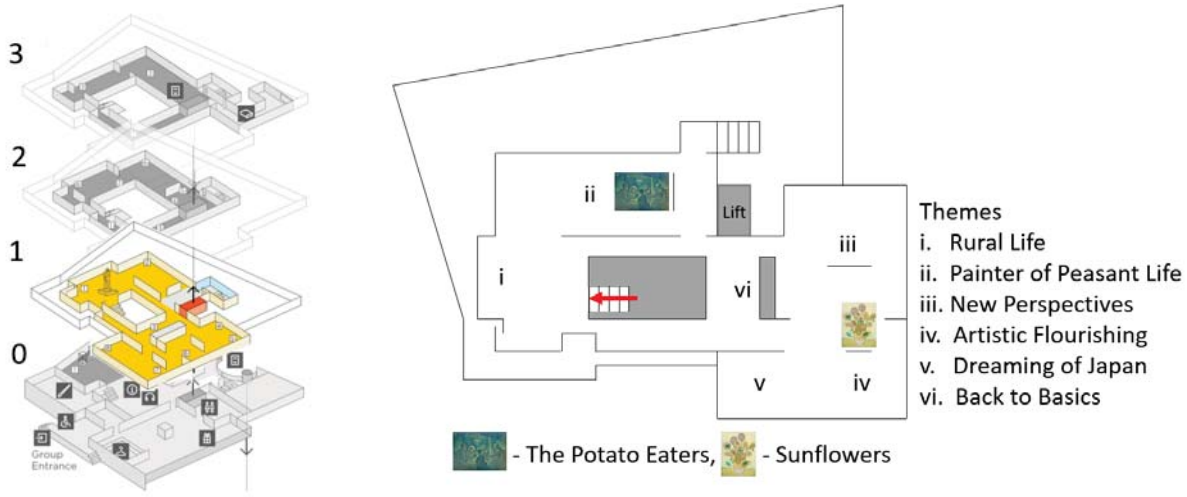

(a) Physical layout across floors.

(b) Layout on Floor 1.

Figure 2: Spatial configuration of the Van Gogh Museum.

We also constitute a record of layout changes that were made during the year 2019 such as temporary exhibitions. As explained subsequently, a fraction of these layout changes result from construction and building maintenance works: as such, they correspond to "exogenous" interventions on the layout. The resulting relocation of artworks is used as natural experiments to validate our predictive model for visitor activity.

3.2 Multimedia tour guide data and pathways

Multimedia tour guide. Our analysis of visitors' paths is based on large-scale datasets describing their activity within the Van Gogh Museum. Specifically, the data was collected via the session logs of MMT guides used by visitors between 2019 and 2021; a few observation records are shown in Table 7, Appendix A. It is worth noting that MMT guides or similar smartphone applications are offered by most major museums around the world. The device has a touchscreen user interface that lets the visitor choose a pre-recorded audio track or visual media content. Each track provides context about a specific artwork or a group of related artworks, which helps better understand its significance in Van Gogh's artistic legacy (Van Gogh Museum 2020). The MMT tracks are available for 45 artworks in 11 languages with a duration for the English language selection ranging from 50 seconds (for "*Peasant Heads*") to 112 seconds (for "*Tree Roots*"). These MMT guides focus on a subset of the artworks on display in the museum at any given time, which is selected by the curators to shed light on Van Gogh's most decisive creative periods. Specifically, the MMT guides display two different tour types, comprising permanent exhibits of the museum collection: (a) the *Highlights* tour covers the masterpieces and most prominent artworks in the collection, and (b) the *Leisure* tour covers a wider range of artworks that includes the aforementioned highlights. Figure 3 shows the different ways in which these options are made available to the visitor - the highlights are displayed in a list format while the leisure tour options are presented in a circular phone-dial format with the artwork identification numbers assigned randomly by the museum.

(a) MMT guide equipment. (b) A list for the Highlights Tour. (c) A wheel for the Leisure Tour.

Figure 3: MMT guide interface.

We leverage the structured data logs generated by visitor sessions on the MMT guides to construct a large-scale panel dataset describing visitors' movement and interactions with the art-

works. In the period considered for our analyses, 25% (25^{th} percentile) to 31% (75^{th} percentile) of all visitors choose to buy the MMT guide service, for a nominal cost in addition to the entry ticket. Past surveys carried out by the museum indicate that visitors with children and those of non-Dutch origin are more likely to buy an MMT guide (2-t t-test p-value < 0.01 for both), but there is no difference in the age groups across the two populations (2-t t-test p-value > 0.60). In the remainder of our analysis, our main dataset records 715,000 distinct visits in 2019.

Clickstream data. We use the clicks on the multimedia guide to determine the progress of the visitor in the museum. This approach has some advantages and disadvantages compared with other location-based technologies such as RFID. On the one hand, the stream of clicks does not reveal the precise location, but it indicates the content viewed by the visitors at different points in time. On the other hand, despite RFID's ability to track exact physical movement in a physical space (Larson et al. 2005, Thiesse and Fleisch 2008), it cannot easily capture the salience of engagement between the customer and the products on offer. For example, in a physical store, the geographical positioning of a customer is "assumed" to be an indicator of her interaction with products in her vicinity - the only foolproof evidence of which being the buying activity recorded at the end of the visit (Hui et al. 2009a). In contrast, clickstream data provide a more accurate signal of user engagement with regard to the surrounding environment. For example, in e-commerce, a customer can only interact with one webpage at each point in time (De et al. 2010, Xu et al. 2014). Hence, in our study, we combine the digital tracking mechanism enabled by the MMT guide with contemporaneous geolocation data for the collection of artworks, to measure visitor engagement with the exhibits displayed by the museum. Finally, we also obtain information on the visitor's origin through the language chosen in the MMT guide. Table 8 in Appendix A shows the breakdown of sessions by language: about a third of visitors choose the English language, while 8% choose Dutch, the local language.

Exploratory analysis. Our modeling approach views visitor activity as a sequence of choices among a finite set of alternatives. That is, each visitor is a decision-maker, who, starting from a given position within the spatial configuration of the museum, chooses in each step the next artwork to transition to, or whether to end the visit altogether. Appendix B illustrates one such sequence undertaken by a randomly selected visitor on the first floor of the museum. By geolocating the artworks in the museum, we are able to estimate the distance traversed for each transition. In our data, we see that a visitor views an artwork only once during her visit, which reflects that there is no value in repetition (in contrast with other content types such as music which users can listen to repeatedly). We use this observation to support one of our model

assumptions, namely that an artwork already seen in the past is excluded from the choice set available for future transitions. Barring this lack of repetition, visitors freely move in the museum space: they begin their visit conventionally on the ground floor, and progressively move upwards (the more common path) or take the lift to the third floor and move downwards; they can come back to the floors that they have already visited and always have the option to terminate their visit and leave the museum.

An exploratory analysis reveals that, on average, the museum features 28 artworks on the MMT guide available for viewing, out of which visitors view 15.57 artworks. Figure 4, which depicts the most frequent paths followed by visitors, shows that in June 2019 (before the COVID-19 pandemic with free-flow traffic), the general trend of visitor movement was much different than that in June 2020 (with COVID-19 social distancing rules and one-way traffic).

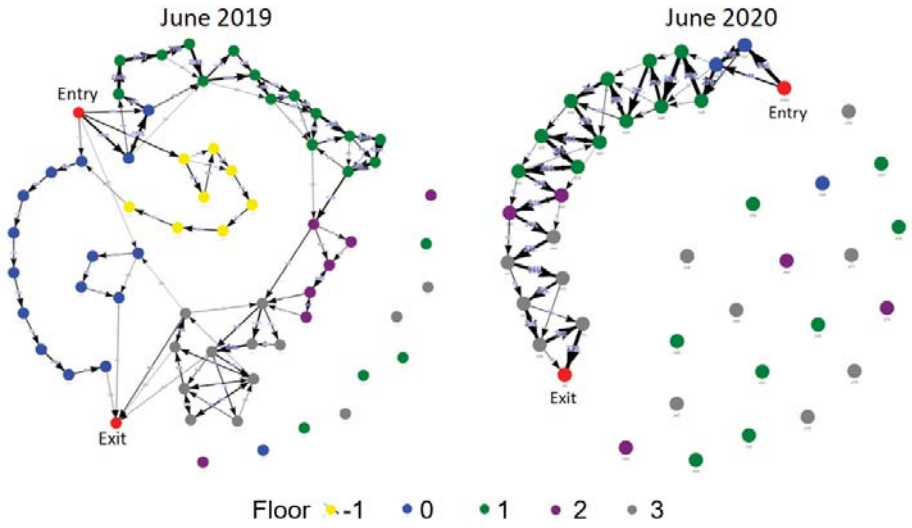

Figure 4: Exploratory transition graphs for visitor movement. Each node represents one artwork and the thickness of the directional arrow between two artworks represents the number of transitions between them. Arrows between only those nodes that saw more than 200 transitions in 2019 and 30 transitions in 2020 are displayed.

3.3 Artistic attributes

WikiArt.org. In addition to internal data provided by the partner museum, we utilize artwork-level data from WikiArt.org. Much like Wikipedia.org, this online repository of information is publicly regulated and it provides a "visual encyclopedia", which is specialized in art-related topics. Detailed information about artworks such as their origin, commission date, style (e.g. for the artwork "Potato Eaters", the style identified is *Realism*), genre (e.g. *Genre-painting*), and the subject matter represented through tags (e.g. *Countryside*, *Mealtimes*) is made available in a standardized format, which can be cross-referenced. We use these external measures to

describe several characteristics of the artworks in the museum collection. Table 9 in Appendix A describes a sample observation from this dataset.

Deep Residual Learning. As an alternative method to track the artistic attributes of artworks in our analyses, we resort to the recent advancements in computer vision models based on *Convolutional Neural Networks* (CNN) in applications involving image/face recognition, object detection, image classification, etc. Specifically, we use a Residual Neural Network (ResNet) developed by He et al. (2016), which was pre-trained on a corpus of over 1.28 million images from ImageNet.org. As a result, it is able to differentiate between a wide range of images irrespective of noise such as lighting, positioning, and contrast. ResNet inspects the pixel-level details of every artwork and provides a 1000-dimensional vectorized representation of each one, providing an objective and scalable encoding of the image features, free from the subjective biases which inevitably arise in expert consultations (Banerjee et al. 2022). See Appendix C for more details.

4 Model for Sequential Transitions

4.1 Probabilistic specification

We start by formulating a general random path choice model, which we call the Pathway MNL model (P-MNL). In particular, we make explicit how visitors' sequential choices are influenced by the layout. While our model description is tailored to the museum context and terminology, it is straightforward to adapt it to other settings; e.g., in retailing, each "visitor" is a "customer", "artworks" correspond to "products", and "exhibition spaces" correspond to "display locations".

Layout decisions. We represent the artworks as a collection $\mathcal{N} = \{1, \ldots, n\}$ of n distinct items. The exhibition space of the museum is captured through n corresponding locations $\mathcal{L} = \{\ell_1, \ldots, \ell_n\}$. In this context, the *layout decision* corresponds to an assignment π that maps each artwork $i \in \mathcal{N}$ to a different location $\pi(i) \in \mathcal{L}$, i.e., $\pi : \mathcal{N} \to \mathcal{L}$ is bijective. For convenience of notation, we denote by 0 the outside option reached at the end of each visit, meaning that 0 captures the visitor's choice of leaving the museum. By convention, the outside option is assigned to a dummy location, denoted by ℓ_0; for all practical purposes, ℓ_0 can be viewed as the museum entry/exit.

Random paths. For a fixed artwork-to-location assignment π, we proceed to describe the random path $\sigma^{v,\pi}$ induced by π for any given visit context v. Here, the index v represents all state variables pertaining to the visitor and the museum that may affect her path, i.e., this

may include visitors' characteristics (e.g., language), museum-level attributes (e.g., congestion), and any dynamic state variable representing the visitor's path so far (e.g., path length so far). The random path $\sigma^{v,\pi} = (\sigma_1^{v,\pi}, \ldots, \sigma_{\ell^{v,\pi}}^{v,\pi})$ is an ordered subset of $\mathcal{N}$ terminating with item $\sigma_{\ell^{v,\pi}}^{v,\pi} = 0$, which describes the sequence of artworks chosen by the visitor until reaching the outside option. Here, we use $\ell^{v,\pi}$ to denote the total number of items and $\sigma_k^{v,\pi}$ is the random variable corresponding to the k-th artwork on her path for every $k = 1, \ldots, \ell^{v,\pi}$. For simplicity of notation, we will omit the reference to the superscripts v, π whenever the context and assignment are fixed.

Pathway MNL model. As common in the choice modeling literature, we impose a probabilistic structure on the distribution of the random paths $\sigma^{v,\pi}$. We assume that $\sigma^{v,\pi}$ is a Markov chain formed by sequential random choices, each of which follows an independent conditional logit experiment. We refer to this probabilistic structure as the *Pathway MNL* model. Specifically, for each transition from an artwork i located in ℓ_i to an artwork j located in ℓ_j, our utility function is additive with respect to two factors: (i) a location-dependent nominal utility $u_l(\ell_i, \ell_j|v)$ associated with the artworks' locations (ℓ_i, ℓ_j), and (ii) an artwork-dependent nominal utility $u_a(i, j|v)$ associated with the pair of artworks (i, j). Both utility functions depend on the context v, which includes museum-level and visitor-level state variables. Consequently, we generate $\sigma^{v,\pi}$ according to the following stochastic process:

- *Initial artwork ($k = 1$):* Each $i \in \mathcal{N}$ is chosen as the initial artwork with probability:

$$\Pr\left[\sigma_1^{v,\pi} = i\right] = \frac{e^{u_l(\ell_0,\pi(i)|v)+u_a(0,i|v)}}{\sum_{j\in\mathcal{N}} e^{u_l(\ell_0,\pi(j)|v)+u_a(0,j|v)}} \ .$$

- *Artwork transitions ($k \geq 2$):* Having sampled the $k-1$ first artworks $(\sigma_1^{v,\pi}, \ldots, \sigma_{k-1}^{v,\pi}) = (s_1, \ldots, s_{k-1})$, let $S_k = \mathcal{N} \setminus \{s_1, \ldots, s_{k-1}\}$ denote the subset of remaining items, including the outside option. Each artwork $i \in S_k \cup 0$ is chosen next with probability:

$$\Pr\left[\sigma_k^{v,\pi} = i \middle| \sigma_1^{v,\pi} = s_1, \ldots, \sigma_{k-1}^{v,\pi} = s_{k-1}\right] = \frac{e^{u_l(\pi(s_{k-1}),\pi(i)|v)+u_a(s_{k-1},i|v)}}{\sum_{j\in S_k} e^{u_l(\pi(s_{k-1}),\pi(j)|v)+u_a(s_{k-1},j|v)}} \ .$$

Parametric form and identifiability. The aforementioned stochastic process can be viewed as a dynamic choice model with two structural assumptions: *(1) Nested choice sets:* Visitors do not see the same artwork more than once; as mentioned earlier, this assumption is in line with our data as well as other related studies (Bourdeau and Chebat 2001). Thus, each state is described by the system state v, the current artwork, and the subset of remaining artworks. Note that it is possible to allow multiple repetitions of the same option, by including a disutility term that

increases in the number of times the option has been consumed already. *(2) Myopic MNL choices:* Visitors myopically maximize their random utility in every transition. In each step, moving from i (current item) to j (potential target) yields a random utility $u_l(\pi(i), \pi(j)|v) + u_a(i, j|v) + \epsilon_{i,j}$, where $\epsilon_{i,j}$ is an independently sampled standard Gumbel random variable.

Here, the first term $u_l(\pi(i), \pi(j)|v)$ depends on the respective locations assigned to the artworks i and j, but it can also be influenced by the context v. This term can be understood as the disutility created by the layout distances from $\pi(i)$ to $\pi(j)$. The second term $u_a(i, j|v)$ depends on the artworks i, j, again possibly influenced by v. It should be understood as the contributions of artworks i and j to the utility. For example, this function may include artwork fixed effects for i and j: one might expect that the propensity of transitioning towards j depends on how attractive j is and whether the experience of i keeps the viewer engaged.

For purposes of normalization, we impose without loss of generality that $u_a(i, 0|v) = u_l(\ell, 0|v) = 0$ for all $i \in \mathcal{N}$ and $\ell \in \mathcal{L}$. Naturally, further parametric assumptions are necessary in order to identify and estimate the utility functions $u_l(\cdot)$ and $u_a(\cdot)$ from realistic data sources. For example, it is not possible to simultaneously identify artwork and location fixed effects, respectively captured in $u_l(\cdot)$ and $u_a(\cdot)$. Indeed, one challenge inherent to museum data is that the layout rarely varies: most artworks remain at fixed locations. In this context, it is only possible to disentangle the location effects from the artwork effects insofar as we adopt a suitable parametrization of the nominal utilities and exploit the variation created by the visitors' sequential movement. In this spirit, we formulate a number of hypotheses regarding visitors' sequential choice, which motivate our econometric specification of $u_l(\cdot)$ and $u_a(\cdot)$ in Section 4.3.

4.2 Hypotheses development

Our goal is to express the nominal utilities $u_l(\ell_i, \ell_j|v)$ and $u_a(i, j|v)$ as a function of underlying features describing the pairs of layout locations $(\ell_i, \ell_j) \in \mathcal{L}$ and artworks $(i, j) \in \mathcal{N}$, forming each potential transition in the P-MNL model, along with the contextual factors v. In what follows, we present well-established behavioral constructs, which help us postulate the effects of various factors on the transition utilities.

Layout distance effects. The accessibility of the choice options available in each transition is an important factor of the layout-dependent utility. While the movement across rooms and floors is unconstrained, we do not expect visitors' trajectories to be arbitrary. An artwork in close proximity is more likely to be considered, whereas one that is farther away from the visitors' current position incurs a higher cost of effort (Pancras et al. 2012, Bell et al. 2018, Kabra et al. 2020). Moreover, skipping intermediary artworks to view a distant one may create a feeling

of loss for visitors. For both tourists and locals who have dedicated time to visit a museum, this skipping behavior is associated with an incomplete experience, leading to a subpar utility (Weingarten et al. 2019).

While the distance metric in the physical space plays an influential role in the transition utilities of the visitors, it is also highly pertinent in the digital space (Ursu 2018). In our context, a visitor not only progresses through the spatial configuration of the museum, but also browses through the user interface of the MMT guide. Similar to physical distances, we expect the farther away the potential option is made available on the MMT guide, the lower the utility associated with choosing it (Ghose et al. 2014). We expect that this effect holds true both when a visitor has to switch between tour types, in which case an additional search effort is required, and when she navigates to different artworks within the same tour type. Consequently, we formulate the following hypothesis regarding the effects of distances on the utility term $u_l(\cdot)$.

Hypothesis 1. *The increase in distance in the physical or digital layout between the current location and that of the potential target artwork decreases the utility of the corresponding transition.*

Artistic distance effects. Several dimensions of the artworks such as their subject matter, period in art history, and materiality (e.g., size of the canvas), together build their distinct identities (Bourdieu 1986). Curators invest significant resources in designing layouts that contextualize the artworks and facilitate the transitions between the exhibits viewed by visitors (Falk et al. 1985). These design choices are consistent with the psychological construct of *habit formation*, which explains that viewers experience the aesthetic pleasure from cultural objects that engender a higher *processing fluency* (Reber et al. 2004). This theory implies that visitors prefer to engage with similar artworks consecutively (Wathieu 1997, Baucells and Sarin 2010). Hence, in addition to the distances associated with the layout, our artwork-dependent utility function $u_a(\cdot)$ needs to incorporate the notion of artistic distances between the current artwork and the potential target, which might influence the visitors' choices. On the other hand, the experience of contrasts might also be desirable to visitors who prefer to engage with a new type of artwork at every turn (Ratner et al. 1999, Baucells and Sarin 2013). In this context, a *variety-seeking* behavior implies that higher utilities are associated with transitions between more dissimilar artworks. This leads us to consider the two following opposing hypotheses.

Hypothesis 2(a). *The increase in artistic distance between the current artwork and the potential target artwork decreases the utility of the corresponding transition.*

Hypothesis 2(b). *The increase in artistic distance between the current artwork and the potential target artwork increases the utility of the corresponding transition.*

Time pressure effect. Visitors tend to plan their visits to cultural institutions by setting a certain time budget (Thaler 1999). Inside a museum, as visitors view more artworks, they may experience a form of *time pressure* due to the depletion of their allotted time budget, resulting in disutility. This well-documented phenomenon is also known as *museum fatigue* (Falk et al. 1985, King and O'boyle 2011). Such pressure or fatigue often leads them to opt for ending their visit sooner (Suri and Monroe 2003). Under such conditions, Hui et al. (2009a) show that visitors adapt by changing their transition strategies. Specifically, as more time has been spent in the museum, we expect visitors to be more purposeful and less exploratory in their path (Yoshimura et al. 2014). At this juncture, the visitor assesses the choice set of artworks that are still left to be seen and then seeks to view the exhibits that maximize her remaining utility before she exhausts her budget. Hence, as time pressure builds up, we expect that visitors are more likely to hastily view prominent artworks, those considered masterpieces, and terminate their visit.

Hypothesis 3. *As the cumulative time spent in the museum increases, the utility of transitioning to an artwork decreases, relative to that of the outside option.*

Hypothesis 4. *As the cumulative time spent in the museum increases, the utility of transitioning to an artwork decreases less for prominent artworks.*

Congestion effect. While moving within the museum, visitors encounter other visitors which impacts their experience. On the one hand, Swann Jr and Pelham (2002), Veeraraghavan and Debo (2009) and Debo et al. (2012) show that people strategically select into social environments that provide self-confirmatory feedback. This implies that congestion around a particular artwork makes the visitor gain higher utility by transitioning to the very same exhibit, with the attention it garners acting as a proxy for its quality (Müller-Trede et al. 2018). On the other hand, what we expect to see in our context is the prevalence of the crowding-out effect, which suggests that the presence of a large number of visitors around an artwork would cause the focal visitor to have a subpar utility of viewing the exhibit (Dargnies 2012, Baek and Shore 2020). The behavioral constraint theory suggests that under such circumstances, the focal visitor may perceive the crowding to be intrusive and restrictive, giving rise to avoidance behavior (Stokols 1976, Lu et al. 2013). This demotivates the focal visitor from transitioning to the artworks with higher congestion rates.

Crowding is often associated with an exit-incentivising effect (Milgram 1970, Dargnies 2012, Lu et al. 2013, Batt and Terwiesch 2015). However, in our context, visitors have already paid the museum ticket fee, and thus, they might not renege due to congestion. Specifically, if the underlying attraction for the artwork is higher than the cost associated with covering a large distance, the visitor can skip the most attractive artworks during congestion time and bide her

time by viewing less crowded exhibits before returning to the masterpieces (Kostami and Ward 2009). Hence, the visitor takes detours from the conventional linear path. Under this theory, the total number of artworks in the visitor's path increases, meaning that she is less likely to end the visit at any given time in her path. Accordingly, we postulate the following hypotheses.

Hypothesis 5. *As the congestion in the museum increases, the utility of transitioning to an artwork increases, relative to that of the outside option.*

Hypothesis 6. *As the congestion in the museum increases, the utility of transitioning to an artwork increases less for prominent artworks.*

4.3 Econometric specification

To reflect on the hypotheses formulated in Section 4.2, we express the P-MNL utilities as a function of a rich set of independent variables, constructed by combining our data sources (i.e., clickstream data, layout information, and art-related repository). The primary variables of interest are the ones that affect the utilities associated with the layout distances between pairs of locations, $u_l(\pi(i), \pi(j)|v)$, and the artwork-related effects for pairs of artworks, $u_a(i, j|v)$. Finally, we incorporate contextual independent variables v capturing congestion and time pressure along with several control variables as in Hui et al. (2009a).

Formally, we fix a visitor context v and consider the $(k-1)^{\text{th}}$ artwork $i = \sigma_{k-1}^{\pi,v}$ in her path at location $\ell_i = \pi(i)$. For every potential target artwork $j \in S_k$ located at $\ell_j = \pi(j)$, the transition utility U_{ij}^v from i to j is expressed as follows:

$$
\begin{aligned}
&U_{ij}^v \\
&\left.\begin{aligned}
&= \beta_1 PhysicalDistance_{\ell_i\ell_j}^v + \beta_2 RoomChange_{\ell_i\ell_j}^v \\
&+\beta_3 FloorChange_{\ell_i\ell_j}^v + \beta_4 IsNearestPhysicalDistance_{\ell_i\ell_j}^v \\
&+\beta_5 ReturnedToFloor_{\ell_i\ell_j}^v + \gamma_1 TourChange_{\ell_j}^v + \gamma_2 MMTguideDistance_{\ell_i\ell_j}^v
\end{aligned}\right\} u_l(\ell_i, \ell_j|v) \\
&\left.\begin{aligned}
&+\delta_1 ArtYearDistance_{ij}^v + \delta_2 ArtSizeDistance_{ij}^v \\
&+\delta_3 ArtTagsDistance_{ij}^v + FE_i^{origin} + FE_j^{destination} \\
&+\lambda_1 ArtSeen^v + \lambda_2 ArtSeen^v \times IsHighlight_j^v \\
&+\mu_1 LogCongestion^v + \mu_2 LogCongestion^v \times IsHighlight_j^v \\
&+FE_{Language^v} + FE_{DayOfTheWeek^v} + FE_{Hour^v} + \nu Thoroughness^v + \epsilon_{ij}^v \ .
\end{aligned}\right\} u_a(i, j|v)
\end{aligned}
\tag{1}
$$

In the above specification (1), we first introduce independent variables capturing the disutility terms associated with the distances between locations. Focusing on physical distances, we construct a distance metric that follows a hierarchically separable structure:

(a) $PhysicalDistance^{v}_{\ell_i\ell_j}$: Distance in meters between locations ℓ_i and ℓ_j. If they are on different floors, the distance includes the distance to the stairs.

(b) $RoomChange^{v}_{\ell_i\ell_j}$: A binary variable with value 1 if ℓ_i and ℓ_j are located in different rooms (arranged to reflect different themes and periods of Van Gogh's life) and 0 otherwise.

(c) $FloorChange^{v}_{\ell_i\ell_j}$: A binary variable with value 1 if ℓ_i and ℓ_j correspond to different floors and 0 otherwise.

This hierarchical approach separates the effects of floor changes, room changes, and absolute distances between locations, which enables a fine-grained study of the relationships between the utility and the layout structure. In addition, we use the binary variable $IsNearestPhysicalDistance^{v}_{\ell_i\ell_j}$ taking the value 1 if the artwork j is the closest to location ℓ_i, among all options $j \in S_k$, and 0 otherwise. To control for the fact that visitors might be less likely to return back to a floor they have already been to, we use the binary variable $ReturnedToFloor^{v}_{\ell_i\ell_j}$, where $ReturnedToFloor^{v}_{\ell_i\ell_j} = 1$ if and only if ℓ_j is on a different floor than ℓ_i and the floor of ℓ_j was previously visited by the visitor.

Next, to account for the digital layout, we make use of the specifics of the differential interface of Highlights and Leisure tours as explained in Section 3.2. For capturing frictions in the MMT guide search process, we have $TourChange^{v}_{\ell_j}$ as a binary variable with value 1 if and only if the visitor v is currently on the Highlights tour while the potential target j is exclusively available on Leisure tour. In other words, this independent variable penalizes transitions that require switching the tour type on the MMT guide. Finally, $MMTguideDistance^{v}_{\ell_i\ell_j}$ is the distance between the selection of artworks i and j on either of the tours, i.e., the number of items to cross when scrolling down or turning the wheel on the MMT guide interface. If the visitor needs to switch the tour type, the latter variable is equal to 0 in the same spirit as the above-described hierarchical distance metric.

We now move on to the artwork-related utility terms in Equation (1). In this context, $ArtYearDistance^{v}_{ij}$ provides the absolute difference in years of creation between artworks i and j. In addition to the chronological differences, $ArtSizeDistance^{v}_{ij}$ captures size-related variations: it is defined as the absolute difference in the log-transformed surface areas of artworks i and j measured in cm^2. Moreover, $ArtTagsDistance^{v}_{ij}$ is computed as the opposite of the number of common artistic tags shared by artworks i and j on the publicly regulated online repository of Wikiart.org (see Section 3.3). As an alternative to this covariate, we make use of the artistic attributes obtained through deep learning techniques (CNN) to estimate the artistic distance between a pair of artworks through $ArtDistance_{ij}$ shown in Appendix C. Lastly, we

incorporate fixed effects for i and j, which capture their marginal utility contributions relative to the outside option. The fixed effect of the current artwork i (origin) can be interpreted as its 'engageability' (Besbes et al. 2016) since it affects the likelihood that the visitor's path does not terminate right after i. The fixed effect of the potential target j (destination) can be thought of as its 'clickability': it captures the attractiveness of j relative to other choice options in S_k.

Finally, we incorporate in Equation (1) several independent variables describing the context v of the visit. Since the total utility of the outside option is fixed to 0, these context variables are implicitly interacted with the binary indicator that the potential target j is an artwork, in opposition to the outside option. This explains why the corresponding additive terms are part of the artwork-related utility function $u_a(i, j|v)$. $ArtSeen^{vk}$ is the number of artworks already seen by the visitor at the time of the transition (alternatively, we consider cumulative duration in the robustness check in Table 14 in Appendix D.2), while $LogCongestion^v$ tracks the total number of visitors that begin their visit to the museum together with visitor v in log-transformed and mean-centered format. As an indicator of which artworks are perceived to be masterpieces of the collection, we rely on the listings of the Highlights and Leisure tours from September and October 2019 (the same period in which our model is estimated); the selection is decided by curators, without precise knowledge of visitor behavior, so that it can be considered exogenous in our analysis. Namely, $IsHighlight_j$ is equal to 1 if artwork j is included in the Highlights tour, and 0 otherwise. This independent variable is interacted with $ArtSeen^v$ and then with $LogCongestion^v$ to see how the effects of these contextual factors are moderated by the status of an artwork as a masterpiece. Finally, we incorporate control variables accounting for visitor heterogeneity. The $Language^v$ fixed effects represent the language selected on the MMT guide. The $DayOfTheWeek^v$ fixed effects control for variations across days of the week, and $Hour^v$ controls for within-day variations. To control for idiosyncratic heterogeneity in visitor engagement, we define $Thoroughness^v$ as the average proportion of the MMT guide tracks that was listened to by the visitor until that moment.

5 Estimation Results and Model Validation

In this section, we describe our estimation results and provide extensive validations of our fitted P-MNL model, ranging from out-of-sample predictions to natural experiments.

5.1 Results

We conduct multinomial logistic regressions in accordance with the econometric specification of the P-MNL model of Section 4.3. Specifically, we construct a dataset of visitor choices based on

the trajectories of 25,000 randomly selected visitors during the months of September and October 2019. Note that the exact layout and arrangement of artworks were tracked by the museum and remained stable throughout this period. We have replicated our estimation analysis on four other time periods of 2019 and obtained similar findings. Table 1 summarizes our estimation results for several variants of the P-MNL model (see also Table 12 in Appendix D.2).

Model X provides an extended specification of the transition utilities that includes fixed effects for all pairs of artworks $(i, j) \in \mathcal{N}$, in replacement of the layout and artistic factors. It is worth noting that Model X is not practically relevant for layout design. Indeed, since its parametric form no longer expresses the transition utilities as a function of underlying features associated with the layout, Model X does not describe the relationships between physical distances and transition utilities, and thus, it cannot generalize to different layouts. Despite also being computationally taxing, it serves as a benchmark to assess the predictive performance of our main models.

We now turn our attention to models that are based on the specification in Equation (1) with increasing sophistication. We see in the progression from Model 0 to Model 4 that the additional factors being considered provide a better fit reflected in the Normalized Log-likelihoods (LLH). It is worth noting that the physical and digital layout-related factors contribute the most towards explaining the variance in the observed transitions – the LLH improves the most for Model 1 and 2 relative to Model 0 (25% and 40%, respectively). Our estimates of the model parameters remain consistent across the different specifications, which indicates their robustness. The fixed effects $Artwork_i$ and $Artwork_j$, corresponding to the origin ("engageability") and destination ("clickability") artworks in the transition, are reported in Table 11 in Appendix D.1.

We now discuss the relationships captured by Model 3 in greater detail. In particular, we calculate how each independent variable affects the attractiveness level $e^{\bar{U}_{ij}^v}$ of the corresponding choice option j, where $\bar{U}_{ij}^v$ is the deterministic portion of the utility in Equation (1). In our numerical application, the transition probabilities tend to be small. Thus, in this regime, the percentage change in the attractiveness level is a good proxy for the change in the corresponding choice probabilities.

	$Transition_{ij}^{v}$					
	(X)	(0)	(1)	(2)	(3)	(4)
Physical Layout Factors						
$IsNearestPhysicalDistance_{\ell_i\ell_j}^{v}$			3.397***	0.583***	0.603***	0.395***
			(0.006)	(0.008)	(0.005)	(0.004)
$PhysicalDistance_{\ell_i\ell_j}^{v}$				-0.287***	-0.282***	-0.296***
				(0.002)	(0.000)	(0.000)
$RoomChange_{\ell_i\ell_j}^{v}$				-2.318***	-2.400***	-2.386***
				(0.030)	(0.006)	(0.006)
$FloorChange_{\ell_i\ell_j}^{v}$				-6.836***	-7.028***	-6.123***
				(0.035)	(0.005)	(0.005)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times RoomChange_{\ell_i\ell_j}^{v}$				0.236***	0.232***	0.248***
				(0.002)	(0.000)	(0.000)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times FloorChange_{\ell_i\ell_j}^{v}$				0.314***	0.310***	0.308***
				(0.002)	(0.000)	(0.000)
$ReturnedToFloor_{\ell_i\ell_j}^{v}$				-0.855***	-0.653***	-0.727***
				(0.014)	(0.012)	(0.011)
Digital Layout Factors						
$TourChange_{\ell_j}^{v}$				-3.093***	-2.993***	-2.104***
				(0.018)	(0.016)	(0.013)
$MMTguideDistance_{\ell_i\ell_j}^{v}$				-0.012***	-0.014***	-0.014***
				(0.000)	(0.000)	(0.000)
Artistic Factors						
$ArtYearDistance_{ij}^{v}$					0.069***	-0.013***
					(0.002)	(0.002)
$ArtSizeDistance_{ij}^{v}$					0.639***	0.416***
					(0.005)	(0.005)
$ArtTagsDistance_{ij}^{v}$					-0.036***	-0.067***
					(0.003)	(0.003)
Contextual Factors						
$ArtSeen^{v}$					-0.206***	
					(0.001)	
$ArtSeen\_Q1^{v}$						6.080***
						(0.125)
$ArtSeen\_Q2^{v}$						4.401***
						(0.034)
$ArtSeen\_Q3^{v}$						2.903***
						(0.024)
$ArtSeen\_Q4^{v}$						2.624***
						(0.010)
$ArtSeen^{v} \times IsHighlight_{j}^{v}$					0.115***	
					(0.001)	
$LogCongestion^{v}$					0.565***	
					(0.014)	
$LogCongestion\_Q1^{v}$						-0.042*
						(0.018)
$LogCongestion\_Q2^{v}$						0.051**
						(0.017)
$LogCongestion\_Q3^{v}$						0.032
						(0.017)
$LogCongestion\_Q4^{v}$						0.000
						(0.017)
$LogCongestion^{v} \times IsHighlight_{j}^{v}$					-0.240***	
					(0.008)	
$Thoroughness^{v}$					3.096***	3.112***
					(0.011)	(0.010)
Fixed Effects						
$Artwork_i \times Artwork_j$	Yes	No	No	No	No	No
$Artwork_i$	No	Yes	Yes	Yes	Yes	Yes
$Artwork_j$	No	Yes	Yes	Yes	Yes	Yes
$Language^{v}$	Yes	Yes	Yes	Yes	Yes	Yes
$Hour^{v}$	Yes	Yes	Yes	Yes	Yes	Yes
$DayOfTheWeek^{v}$	Yes	Yes	Yes	Yes	Yes	Yes
AIC	1570.93	173.69	176.27	192.61	210.66	216.61
Normalized Log-likelihood	-1.07	-2.30	-1.72	-1.38	-1.33	-1.38
Num. parameters	786	88	89	97	106	109
Num. events	392,209	392,209	392,209	392,209	392,209	392,209
Num. obs.	6,221,248	6,221,248	6,221,248	6,221,248	6,221,248	6,221,248

$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$. Standard errors in parentheses.

Table 1: Pathway MNL estimation results

Layout distance effects. Our hierarchical representation of distances through the absolute distance in meters, room, and floor changes is associated with significant negative coefficients. Specifically, when the current artwork i and the potential target j are located on the same floor and in the same room, as the distance between them increases by 1 meter, visitors are 24.6% ($e^{-0.282} = 0.754$) less likely to transition to artwork j. However, visitors are much less sensitive to absolute distances when there is a room or floor change (variation of $1 - e^{-0.282+0.232} = -4.8\%$ and $1 - e^{-0.282+0.310} = +2.8\%$ per additional meter, respectively), but at the expense of significant drops in transition rates of 91.0% ($e^{-2.400} = 0.090$) and 99.9% ($e^{-7.028} = 0.001$), respectively. In other words, it is unlikely to observe transitions across rooms and floors, as visitors often focus on their options in a short range. This implies that transitions to the next room tend to occur once nearby options have been exhausted. Moreover, as expected, the coefficients of the variables $IsNearestPhysicalDistance_{ij}^{v}$, which tracks the nearest artwork from every artwork i in the choice set, and $ReturnedToFloor_{ij}^{v}$, which indicates that visitors return to a floor they have already been to, are significant with positive and negative effects, respectively. This implies that visitors are 82.7% ($e^{0.603} = 1.827$) more likely to transition to artwork j if it is the nearest option available as also observed by Ahmadi (1997) in the case of amusement parks, and 48.0% ($e^{-0.653} = 0.520$) less likely if it is on a different floor that has been visited before, in line with the findings of Bourdeau and Chebat (2001). In summary, these results are consistent with Hypothesis 1 for all dimensions of our distance metric, thereby demonstrating the strong effects of the physical layout on visitor pathways.

Similarly, our notions of digital layout distance, in the form of search efforts on the MMT guide interface, are also associated with negative effects on the transition probabilities. Once visitors select the Highlights tour, they are 95.0% ($e^{-2.993} = 0.050$) less likely to transition towards artworks exclusively displayed in the Leisure tour. In other words, being on one type of MMT guide tour often makes those on the other selection go unnoticed. Moreover, within the same tour type, the farther away an option is on the interface of the MMT guide, the less likely the transition is to that artwork: the choice probability drops by 1.4% per position to scroll down ($e^{-0.014} = 0.986$). These findings indicate that the MMT guide interface has a significant impact on visitors' activity. Recall that the artwork IDs are randomly assigned by the museum and tour guide provider.

Artistic distance effects. The artistic characteristics of the pair of artworks forming each transition also significantly contribute to explaining the transition choices, although with smaller effects compared to the impact of the physical and digital layouts. On the one hand, the transition probability increases with chronological and size differences between the current artwork

and the potential target. This relationship is consistent with a variety-seeking behavior, as formulated in Hypothesis 2(b). On the other hand, a unit increase in the dissimilarity of the subject/theme represented in the artworks comes with a decrease of the corresponding transition probability. This finding is corroborated with the analyses carried out by using $ArtDistance_{ij}^{v}$ as the metric of artistic distance based on deep learning techniques (CNN) discussed in Section 3.3 (see results in Table 13). Accordingly, the behavior exhibited by visitors is in line with habit formation in Hypothesis 2(a). To refine our initial hypotheses, our results suggest that visitors might prefer variations in the art's material qualities, but might favour consistency in the artistic subjects. It would be interesting for future research to study and validate such hypotheses.

Time pressure effect. We now consider the effect on the utility of the variable $ArtSeen^{v}$, which measures the cumulative number of artworks already visited. Every additional artwork reduces the propensity to visit another one by 18.6% ($e^{-0.206} = 0.814$), relative to the outside option. Hence, this relationship is in line with the time pressure phenomenon described in Hypothesis 3. Furthermore, Model 4 shows that this effect is very strong for the initial transitions, and less pronounced afterwards. This is because moving from quartile 1 (Q1) to quartile 2 (Q2) results in a decrease in the transition probability to an artwork in comparison to the outside option by 81.3% ($e^{4.401-6.080} = 0.186$), while moving from Q2 to Q3 and from Q3 to Q4 to a decrease of 77.6% and 24.4% respectively ($e^{+2.903-4.401} = 0.223$ and $e^{+2.624-2.903} = 0.756$). In addition, the time pressure factor is still negative but of lesser magnitude for the masterpieces, as indicated by the significant positive coefficient +0.115 of the corresponding interaction term. As a number of artworks seen increases, the relative probability of transitioning to the main highlights increases compared to that of less prominent artworks. This observation is consistent with Hypothesis 4: visitors tend to focus on the masterpieces of the collection at the end of their visit.

Congestion effect. Finally, we examine how congestion within the museum affects visitors' choices. We find that the effect of $LogCongestion^{v}$ is positive and significant, which is consistent with Hypothesis 5. Specifically, an increase of congestion by 10% implies an increase of the transition probabilities towards artworks by 5.5% ($1.1^{0.565} = 1.055$), relative to the outside option. In congested environments, visitors are more likely to continue engaging with artworks on the MMT guide, as opposed to terminating their path. Model 4 shows that this effect is concave increasing: when congestion is very low, moving from Q1 to Q2 in congestion results in a 10.3% increase in the probability of continuing the visit ($e^{0.051+0.042} = 1.103$). As the congestion increases from Q2 to Q3 and from Q3 to Q4, this probability slightly reduces but

remains statistically insignificant. Finally, the effect of congestion is negatively and significantly moderated by the artwork prominence, with a significant coefficient of -0.240. This provides evidence for the hypothesized changes in visitor preferences outlined in Hypothesis 6. Indeed, an increase in congestion results in the concentration of crowds around the popular highlights, nudging the focal visitor to transition to artworks that the museum signals to be less prominent. Taken together, these findings corroborate the initial intuition resulting from the discussion between the authors and museum managers about crowd behavior: the visitor traffic is self-regulating with higher levels of congestion forcing visitors to spread out and view artworks they would not have seen otherwise.

5.2 Out-of-sample validation

We proceed to evaluate the predictive performance of the fitted P-MNL model on new unseen data, and thus, provide external validity to this modelling approach. For this, we construct an out-of-sample test set by taking another subsample of 25,000 visitors in the same time span as our training set. The quality of the predictions of our main model (Model 3) is then compared with that of several benchmarks using Normalized Log-Likelihood (LLH), classification accuracy, and Mean Absolute Percentage Error (MAPE).

Predictive power. For this comparison, we consider both simplified and extended specifications of the P-MNL model. More specifically, to quantify the gains in prediction accuracy from layout-related factors, we use Models 0 and 1 from Table 1 as simplified benchmarks. Regarding the more complex benchmarks, we test out the extended specification of Model X, which includes fixed effects for all pairs of artworks. Recall that this model cannot generalize to different layouts, but it arguably offers the best-possible predictive power within our family of nested P-MNL models. Additionally, we construct Model 18 (Table 15, Appendix D.2) as a variant of our main model in which we replace the binary variables $FloorChange^{v}_{\ell_i,\ell_j}$ with floor-to-floor fixed effects representing each potential transition amongst floors $\{E, 0, 1, 2, 3, T\}$, where E stands for the entry and T for the exit.

Table 2 displays our predictive performance metrics across the different specifications of the Pathway MNL model (also see Table 16). Comparing Models 0 and 1, the incorporation of the nearest variable indicators suffices to double the classification accuracy to 48.6%. Considering that each visitor faces 11 choice options in each transition on average, this level of accuracy is a striking indication of the strong layout effects. Although this naive benchmark achieves a satisfactory level of accuracy, our main model (Model 3) leads to a further 22% increase in LLH, a 29% increase in classification accuracy and a 27% reduction in MAPE relatively to Model 1.

As we consider more complex P-MNL specifications, the improvements in prediction accuracy are more marginal. Model 18 provides a 5.3% increase in LLH, a 1.6% increase in classification accuracy, and a 6.8% reduction in MAPE relatively to our main model. As expected, Model X achieves the highest predictive performance. The gap conceded by Model 18 against Model X shows the cost of our parametric form on the transition utilities. Yet, as explained earlier, this gain in predictive accuracy comes at the cost of computational complexity and the inability to generalize to modified layouts, which differ from the one in operation in October 2019. In this context, the advantages of our main model will be further illustrated in Section 5.3.

Model	Reference	Log-likelihood	Classification accuracy
Without layout factors	Table 1 Model 0	-2.31	23.45%
With the nearest artwork	Table 1 Model 1	-1.72	48.68%
Main model	Table 1 Model 3	-1.34	63.00%
With floor to floor FE	Appx D.2 Table 15 Model 18	-1.27	64.06%
With pair-level FE	Table 1 Model X	-1.08	67.20%

Table 2: Out-of-sample predictive performance of P-MNL models. Note: Log-likelihood is normalized, i.e., we report the ratio of log-likelihood with the number of observations, and hence this figure reports the average value of the log-probability of the actual choice in the model.

The value of a sequential model. Although our previous analysis provides an out-of-sample validation for the marginal choice probabilities prescribed by the MNL model in each transition, it does not inform us about the validity of the distribution over paths $\sigma^{v,\pi}$ generated by the P-MNL model. In particular, our key performance indicators for visitor engagement are the *hit rate* of any item $i \in \mathcal{N}$, defined as the probability $\Pr[i \in \sigma^{v,\pi}]$ that a visitor reaches item i, and the *total number of hits* is defined as the expected total number of items reached by the visitor $\mathbb{E}[|\sigma^{v,\pi}|] = \sum_{i \in \mathcal{N}} \Pr[i \in \sigma^{v,\pi}]$. With respect to the P-MNL model, these probabilities are estimated using Monte Carlo simulation. As a benchmark for predictive performance, we consider a class of logistic regression models, denoted by LOGIT, which is separately fitted for each item $i \in \mathcal{N}$ (see Appendix D.4). We mention in passing that this class of models is subject to the same limitation as Model X: it cannot generalize to modified layouts. For a finer comparison between P-MNL and LOGIT, we estimate the hit rates and the total number of hits conditional to the first k elements in the visitor path $\sigma^{v,\pi}_{\leq k} = (\sigma^{v,\pi}_1, \ldots, \sigma^{v,\pi}_k)$ where the parameter k is varied in $\{0, 5, 10, 15\}$. In the case of LOGIT, we need to make a further assumption of probabilistic independence to estimate the joint distribution of hits.

The predictive performance of P-MNL and LOGIT is reported in Table 3; each metric is averaged over all visitors and/or artworks in the test set. As one might expect, our model concedes a gap of 6% in LLH for $k = 0$. This fact is not surprising as the LOGIT benchmark

represents a direct relationship between the visit context v and the hit rate $\Pr[i \in \sigma^{v,\pi}]$. By contrast, this relationship is only indirectly captured by P-MNL via the distribution of sequential choices. Moreover, given that a separate LOGIT is fit for each $i \in \mathcal{N}$, P-MNL has fewer degrees of freedom. However, as k increases, the important observation is that the LLH value for the P-MNL model improves from -0.50 ($k = 0$) to -0.36 ($k = 15$), whereas it degrades from -0.47 ($k = 0$) to -0.87 ($k = 15$) for the LOGIT model. In other words, P-MNL accurately describes the correlations of hits along visitors' paths, whereas the assumption of probabilistic independence is inadequate in the case of the LOGIT model. The linear correlation between the observed and predicted total number of hits corroborates this insight by exhibiting a much steeper increase for the P-MNL in comparison to LOGIT as k increases (approximately 9x increase from $k = 0$ to $k = 15$ for P-MNL versus 3x increase for LOGIT). Figure 5 further shows that the distribution of the total number of hits per visitor under the fitted P-MNL model is much closer to the actual one, compared with the corresponding LOGIT model.

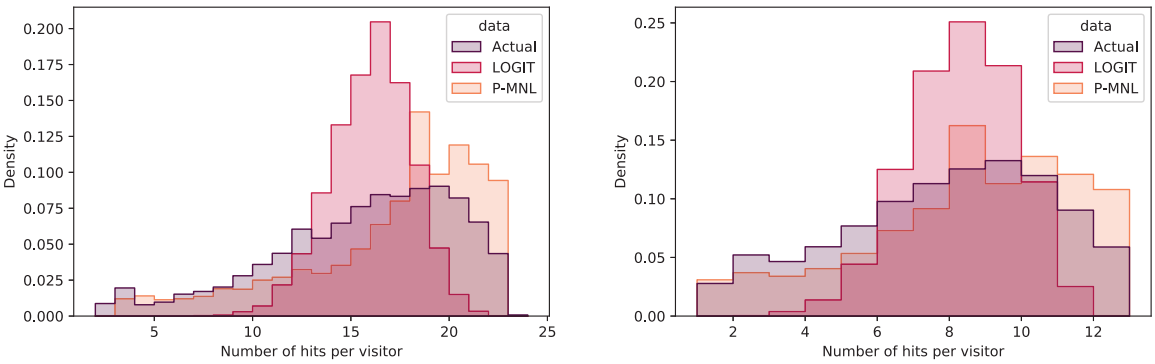

(a) $k = 0$ (b) $k = 10$

Figure 5: Distributions for the total number of hits: actual distribution compared with simulated distribution using P-MNL and LOGIT.

k	Log-likelihood		Correlation	
	P-MNL	LOGIT	P-MNL	LOGIT
0	-0.50	-0.47	0.08	0.13
5	-0.53	-0.56	0.33	0.20
10	-0.47	-0.71	0.50	0.10
15	-0.36	-0.87	0.70	0.32

Table 3: Out-of-sample predictive performance for the P-MNL and LOGIT models as a function of k. LLH is computed with respect to the hit rate per visitor and per artwork and correlation is computed with respect to the total number of hits per visitor.

5.3 Natural experiments

The last step of our model validation is to evaluate its ability to predict how changes to the layout affect the visitor pathways. For this, we leverage four natural experiments caused by construction works and building maintenance operations in certain galleries of the Van Gogh Museum at different points in time in November and December 2019. As result, several artworks were relocated to the ground floor, as described in Table 4. In our context, these interventions correspond to exogenous changes to the physical layout factors: these temporary changes are neither driven by visitor-related objectives, nor pre-announced to incoming visitors. We conduct a bias analysis and compare the differences in transition rates predicted by our P-MNL model to the estimated treatment effects associated with these natural experiments.

Dates in 2019	Artworks relocated	From	To	Artworks removed
Nov. 11-14	"The Yellow House" and "Seascape near Les Saint-Marie-de-la-Mer"	Floor 2	Floor 0	"Gauguin: Van Gogh painting Sunflowers"
Nov. 18-22	"Sunflowers"	Floor 1	Floor 0	
Nov. 25-28	"Almond Blossom" and "Wheatfield with a Reaper"	Floor 3	Floor 0	"Irises" and "Copies after Millet"
Dec. 9-12	"The Potato Eaters"	Floor 1	Floor 0	"The Cottage" and "Peasant Heads"

Table 4: Description of the natural experiments

Throughout the analysis, the outcome variables of interest are the artwork pair-level transition rates aggregated at the daily level. First, we estimate heterogeneous treatment effects for each pair of artworks and each natural experiment. Given that these interventions simultaneously affect all units (i.e., pairs of artworks and visitors), there is no apparent control group. As a result, the underlying assumptions for a Difference-in-Differences estimation are not satisfied. Instead, we exploit the weekly seasonality of the visitor behavior to construct a time series-based estimate of the treatment effect, where the control is formed by a simple auto-regressive model (see Appendix D.5). Formally, we define, $ObservedTreatment_{ijt} = \log(Observed_{ijt}) - \log(Control_{ijt})$, where $Observed_{ijt}$ is the observed transition rate from artwork i to artwork j on day t, out of all the visitors that reached artwork i on day t, and $Control_{ijt}$ is the output of our auto-regressive model, which gives an estimate of the transition rate with respect to the original layout.

Next, using the fitted P-MNL model, we simulate the effects of each intervention by predicting the visitor transition probabilities on the new modified layouts, and compare them to those on the original layout. Here, we define $PredictedWithChange_{ijt}$ as the average transition

rate from i to j for day t obtained via simulation on the modified layout. Similarly, we define $PredictedWithoutChange_{ijt}$ as the average transition rate obtained via simulation on the original layout. Consequently, the predicted treatment effect variable $PredictedTreatment_{ijt} = \log(PredictedWithChange_{ijt}) - \log(PredictedWithoutChange_{ijt})$ measures the effect of the layout change predicted by the P-MNL model.

To test the validity of these predictions, we estimate the following simple OLS regression:

$$ObservedTreatment_{ijt} = \beta \cdot PredictedTreatment_{ijt} + \varepsilon_{ijt} \ .$$

In Table 5, we report our results for each of the four interventions over 3 days pre- and post-intervention periods – except for the intervention at Nov. 25-28 in which case the pre-intervention perdiod cannot be extended beyond 2 days to avoid colliding with the previous intervention (Nov. 22). Moreover, we use a subsample consisting only of transition pairs that include the relocated artwork. Table 19 in Appendix D.5 shows similar results across the full sample.

	$ObservedTreatment_{ijt}$				
	Nov. 11-14	Nov. 18-22	Nov. 25-28	Dec. 9-12	Overall
$PredictedTreatment_{ijt}$	0.939***	0.547***	0.617***	0.657***	0.764***
	(0.048)	(0.136)	(0.061)	(0.092)	(0.035)
Pair-level Control	Yes	Yes	Yes	Yes	Yes
R^2	0.736	0.454	0.611	0.499	0.605
Adj. R^2	0.688	0.308	0.496	0.406	0.525
Num. obs.	609	190	396	296	1491

$^{*}p < 0.1;^{**}\ p < 0.05;^{***}\ p < 0.01$.

Table 5: Relating the predicted and estimated treatment effects for four layout interventions.

The estimated regression coefficient ranges from 0.55 (for the intervention between Nov. 18 and 22) to 0.94 (for the intervention between Nov. 11 and 14) with the average effect across all interventions being 0.77. Hence, the natural experiments provide evidence that our model correctly anticipates how the layout changes affect visitors' sequential choices in the museum. Although a coefficient of one would be ideal, because our model would then provide an unbiased estimate of the treatment effects, our results suggests that the impact of the layout change is even larger in magnitude than that captured by our model.

6 Counterfactuals: Designing Enhanced Museum Layouts

6.1 The layout optimization problem

In this section, we formulate the design of the layout as an optimization problem. Formally, the *layout optimization problem* seeks to determine an artwork-to-location assignment $\pi : \mathcal{N} \to \mathcal{L}$ that maximizes the expected total number of hits $\mathbb{E}_{\sigma \sim \mathcal{D}^v_\pi}[|\sigma|]$ along visitors' paths. Here, $\mathcal{D}^v_\pi$ denotes the distribution over random paths σ drawn from the Markov chain model of Section 4.1.

To simplify the formulation of this optimization problem, we take the contextual information v as fixed and hence drop any reference to it in our notation.

The question of choosing an adequate objective for layout design is in itself an equivocal issue. The main mission of cultural institutions is to create engaging and educating experiences. As such, the total number of hits is a relevant metric for visitor engagement. Additionally, as corroborated by managers and curators of the Van Gogh Museum, more hits indicate a higher usage of the MMT guide, which is positively associated with the quality of visitor experience in qualitative surveys. This objective function also aligns with the goal of exposing the audience to the largest possible variety of Van Gogh's works implying that less renowned artworks also receive significant share of attention. Finally, traffic and crowd management is another crucial objective of the museum. By nudging the visitor attention towards certain artworks, the MMT guide affects visitor movement and the congestion in the galleries. Encouraging visitors to consider a more diverse set of artworks could possibly lead to better balance of visitors' queues around the artworks in opposition to concentrating the visitor attention on the main highlights.

In our simulations, we consider a realistic optimization setting that accounts for heterogeneity in v representing different visit types. In practice, the optimal layout should depend on the variety of visitor profiles and congestion levels. We thus consider a more general formulation of the layout optimization problem where v is viewed as a latent random variable and the objective function $\mathbb{E}_v[\mathbb{E}_{\sigma\sim\mathcal{D}_\pi^v}[|\sigma|]]$ proceeds from a probabilistic mixture over the realizations of v.

6.2 Computing layout policies

We first establish a strong hardness result that indicates the computational complexity of the layout optimization problem even for a fixed visitor context v and for simple parametric forms of the P-MNL model. Specifically, we say that the utility function is linear if the nominal utility of each item pair only depends on the potential target, i.e., $u_a(i,j) = u(j)$ for all $i,j \in \mathcal{N}$. By contrast, we use the notion of interaction effects to refer to the general case where $u_a(i,j)$ depends on both i and j.

Theorem 1. *The layout optimization problem is NP-hard even under a linear utility function. In the presence of interaction effects, the layout optimization problem is NP-hard to approximate within some constant $c > 0$.*

These results suggest that finding optimal solutions to the layout optimization problem might be computationally difficult. Leaving further theoretical analyses for future research, we develop simple and intuitive heuristics that exploit the layout structure.

Index policies. We first consider two simple index policies, denoted by π^{CI} and π^{EI}, where the assignment is determined by ranking items based on their 'clickability' and 'engageability', respectively (Gallego et al. 2020, Besbes et al. 2016). Policy π^{CI} locates more clickable artworks farther from the exit to pull visitors' attention toward remote galleries, thereby generating longer paths. Policy π^{EI} locates more engageable artworks closer to an exit; as these locations are more

"leaky", the intuition is that this assignment will reduce the likelihood that visitors terminate their path. More formally, let $\ell = \ell_i, \ldots, \ell_m$ be a numbering of the locations $\ell \in \mathcal{L}$ by increasing $PhysicalDistance_{\ell,\ell_0}$ from the exit. The clickability policy first orders the artworks $j = 1, \ldots, n$ by increasing FE_j and then assigns $\pi^{CI}(j) = \ell_j$. The engageability policy first orders the artworks $i = 1, \ldots, n$ by decreasing FE_i and then assigns $\pi^{EI}(i) = \ell_i$. To avoid relocating artworks across floors, we separately implement these index policies on each floor.

Local search. Next, we consider policy π^{LS} which improves the layout through greedy improvements, where the locations of artworks are sequentially swapped. Initially, we consider the assignment π^0 corresponding to the current layout. In every stage $t \geq 1$, we consider the candidate assignments $\pi^t_{i \leftrightarrow j}$ obtained from π^{t-1} by exchanging the locations of a pair of distinct items $(i, j) \in \mathcal{N}^2$ on the same floor, i.e., $\pi^t_{i \leftrightarrow j}(k) = \pi^{t-1}(k)$ for all $k \neq i, j$ and $\pi^t_{i \leftrightarrow j}(i) = \pi^{t-1}(j), \pi^t_{i \leftrightarrow j}(j) = \pi^{t-1}(i)$. The next assignment is chosen as $\pi^t = \text{argmax}_{(i,j) \in S_t} \mathbb{E}[|\sigma^{\pi^t_{i \leftrightarrow j}}|]$ by maximizing the expected path length out of all candidate swaps in a set S_t. To ensure that the path length may only increase, we fall back to $\pi^t = \pi^{t-1}$ if none of the swaps improves our objective by more than 0.04. For simplicity, we pick S_t as a singleton uniformly chosen over all possible swaps and run our algorithm for 1000 stages.

Constrained local search. One practical drawback of previously discussed heuristics is that they can result in a completely different arrangement of the artworks, which is incompatible with the theme-based narrative in the form of different rooms curated by the museum. To ensure that the current content of each room is not altered, we impose constraints on the set of candidate swaps S_t considered in each step of the local search. This viewpoint yields a more realistic, room-constrained local search heuristic π^{CS}, which is constructed by considering only local swaps of artworks placed in the same rooms.

6.3 Results and insights

The results of the numerical simulations are reported in Table 6. Each entry is obtained by simulating 50 independently generated paths for a random subsample of 3,000 visitors. We find that significant improvements of visitor engagement on the MMT guide can be generated by new layout policies in our simulated environment. First, we see that the local search policy (π^{LS}) lifts the total number of hits by more than 7%, and the room-constrained local search policy (π^{CS}) shows an improvement exceeding 4%. Hence, significant gains can be achieved by implementing only local changes to the layout, which do not affect the thematic structure of the current arrangement of artworks. Furthermore, we observe that our intuitive index policies π^{EI} and π^{CI} are either on par or slightly better than the baseline. These results suggest that clickability and engageability are relevant indicators to enhance the design of the layout. Note that similar index policies in reverse order yield a comparable performance; one possible explanation is that clickability and engageability are positively correlated across the artworks.

	Policy type	Total number of hits $\mathbb{E}[\lvert\sigma^{\pi}\rvert]$	Improvement over baseline
Baseline	-	15.19	-
Indices	Engageability π^{EI}	15.19	-0.03%
	Clickability π^{CI}	15.57	2.50%
Local search	Best single swap	15.30	0.70%
	Room-constrained π^{CS}	15.91	4.69%
	Unconstrained π^{LS}	16.36	7.65%

Table 6: Simulations of counterfactuals for different layout policies and performance metrics.

To better understand the mechanics of the enhanced layouts, we decompose the effects of a single swap. In Figure 6, we visualize in a diagram the percentage change in the hit rates of artworks located in Floor 1 following a swap between "The Potato Eaters" and "The Zouave", which results in nearly a 0.7% increase of the total hit rate. As one might expect, since the location of the former artwork is more accessible, The Zouave's engagement increases by 11.5% whereas The Potato Eaters' engagement sees a relatively smaller improvement of 1.5%. Furthermore, it is striking to see that our swap has positive spillovers on the vast majority of other artworks, initially located near The Zouave. Several artworks in the new vicinity of The Potato Eaters seem to benefit from the higher clickability of this artwork, which attracts more incoming transitions. At the same time, the higher level of engageability of The Zouave might better retain visitors after the encounter with this artwork, thereby also benefiting The Potato Eaters.

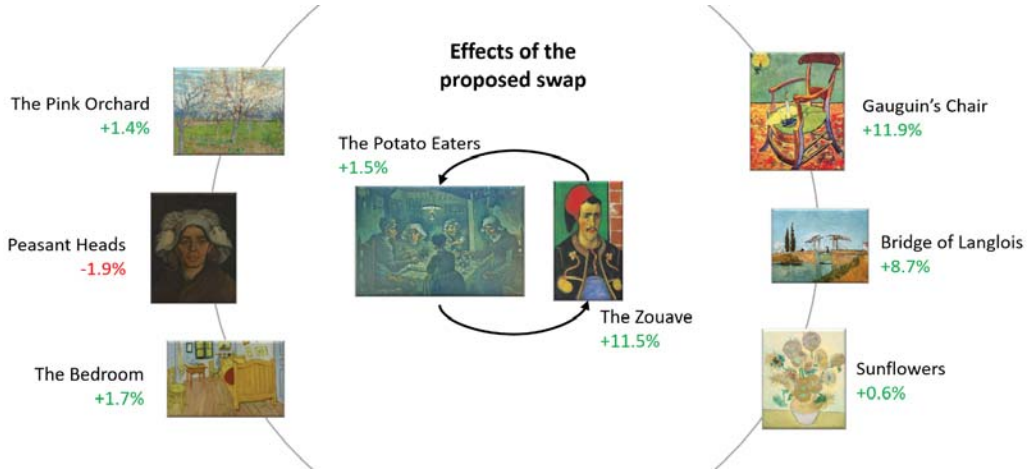

Figure 6: Effect of swapping the position of "The Potato Eaters" with "The Zouave" on the corresponding hit rates for the artworks on Floor 1.

In addition, our results indicate that the proposed policies also generate layouts in which visitors are less concentrated around the highlights in the museum. In other words, the increase in engagement is not driven by higher hit rates for the highlights – those with higher clickability but lower engageability – but the opposite: highlights are moved to less accessible positions causing the less prominent artworks near these locations to benefit from the inflow. The highlights are thus used as "anchors" to generate more activity in relatively less frequented parts of the museum,

and because these locations are further from the exit, they can also improve on retention. (Note that this is somewhat analogous to the rationale for putting milk at the back of the store in retail settings.) Less prominent artworks now take the former positions of the highlights and benefit from increased accessibility. In summary, our policies suggest that engineering layouts using our quantitative approach is a promising direction for improving visitor engagement. Note that our framework can also be easily used to optimize the digital layout of the MMT guide, e.g., the artworks recommended in the Highlights tour. This issue is left for future research.

7 Conclusions

In this paper, we develop a model for sequential customer choices in the context of cultural experiences, *Pathway MNL*, which we validate on large-scale data from the Van Gogh museum. We find that both the physical and digital distances are crucial determinants of the visitor journey, and that different contents and artworks may well vary in their ability to attract and retain visitors in the museum. We illustrate how our model can be employed to optimize the layout and lift visitor engagement on the multimedia guide. This research should be seen as an attempt to support the curatorial decisions of museum managers using a visitor-centric approach. Subtle adjustments to the experiences on offer can contribute to making the museum space more inclusive and to broadening its audience – for example, by driving more visitor engagement using digital channels. Hence, the spirit of our work aligns well with Vincent van Gogh's artistic quest – "I want to touch people with my art," said Vincent in one of his letters to his brother, Theo.

Beyond the application to cultural institutions, our layout optimization framework can be applied to other settings such as retail, entertainment, or education, in which the physical and digital configurations affect the user paths. Hence, this work opens up new opportunities for further research. First, it provides new predictive and prescriptive analytics perspectives for the cultural sector, where big data has not yet generated its full potential and significant practical impact could be achieved (Li et al. 2018). Curating an art collection is a complex endeavour, combining historical, pedagogical, and aesthetic dimensions. In this context, our data-driven simulation tool can support the work of curators and designers for generating and testing out new (efficient and stimulating) display ideas and for planning visitors' trajectories. Second, our analysis underscores the significance of time pressure and congestion on engagement, as in Hui et al. (2009a), so that the planner can modify the experience by introducing opportunities for rest (Baucells and Zhao 2019), e.g., a spacious coffee shop for a pause. Third, it paves the way for field experiments that test different types of interventions, in the physical space (more costly) or in the digital space (cheaper and more flexible), which might uncover new insights. Finally, our framework can inspire new methodological developments for modeling sequential customer choices and more advanced computational analyses for layout optimization problems.

Acknowledgements: We would like to express our sincerest gratitude to Koen Snijders, Rianne van Dam, Laurine van de Wiel, and all the personnel at the Van Gogh Museum who not only supported us in the execution

of this project, but made it an enjoyable and fulfilling experience for us. We also thank the seminar attendees at Indiana University Bloomington, Imperial College London, Indian School of Business, and Georgia Institute of Technology for their valuable comments in shaping this paper.

Victor Martínez-de-Albéniz acknowledges the financial support provided by the Agencia Estatal de Investigación (AEI) from the Spanish Ministry of Science and Innovation, with project reference PID2020-116135GB-I00 MCIN/ AEI / 10.13039/501100011033.

References

Abeliuk A, Berbeglia G, Cebrian M, Van Hentenryck P (2016) Assortment optimization under a multinomial logit model with position bias and social influence. *4OR* 14(1):57–75.

Ahmadi RH (1997) Managing capacity and flow at theme parks. *Operations research* 45(1):1–13.

Aouad A, Segev D (2020) Display optimization for vertically differentiated locations under multinomial logit preferences. *Management Science* .

Baek J, Shore J (2020) Forum size and content contribution per person: a field experiment. *Management Science* 66(12):5906–5924.

Banerjee M, Cole BM, Ingram P (2022) "Distinctive from What? And for Whom?" deep learning-based product distinctiveness, social structure, and third-party certifications. *Academy of Management Journal* (ja).

Batt RJ, Terwiesch C (2015) Waiting patiently: An empirical study of queue abandonment in an emergency department. *Management Science* 61(1):39–59.

Baucells M, Sarin RK (2010) Predicting utility under satiation and habit formation. *Management Science* 56(2):286–301.

Baucells M, Sarin RK (2013) Determinants of experienced utility: Laws and implications. *Decision Analysis* 10(2):135–151.

Baucells M, Zhao L (2019) It is time to get some rest. *Management Science* 65(4):1717–1734.

Bell DR, Gallino S, Moreno A (2018) Offline showrooms in omnichannel retail: Demand and operational benefits. *Management Science* 64(4):1629–1651.

Bernstein F, Martínez-de Albéniz V (2017) Dynamic product rotation in the presence of strategic customers. *Management Science* 63(7):2092–2107.

Besbes O, Gur Y, Zeevi A (2016) Optimization in online content recommendation services: Beyond click-through rates. *Manufacturing & Service Operations Management* 18(1):15–33.

Bourdeau L, Chebat JC (2001) An empirical study of the effects of the design of the display galleries of an art gallery on the movement of visitors. *Museum management and curatorship* 19(1):63–73.

Bourdieu P (1986) The forms of capital .

Brown A, Ratzkin R (2011) Making sense of audience engagement. *The San Francisco Foundation* 1:78.

Centorrino P, Corbetta A, Cristiani E, Onofri E (2021) Managing crowded museums: Visitors flow measurement, analysis, modeling, and optimization. *Journal of Computational Science* 53:101357.

Dargnies MP (2012) Men too sometimes shy away from competition: The case of team competition. *Management Science* 58(11):1982–2000.

Das Gupta A, Karmarkar US, Roels G (2016) The design of experiential services with acclimation and memory decay: Optimal sequence and duration. *Management Science* 62(5):1278–1296.

De P, Hu Y, Rahman MS (2010) Technology usage and online sales: An empirical study. *Management Science* 56(11):1930–1945.

Debo LG, Parlour C, Rajan U (2012) Signaling quality via queues. *Management Science* 58(5):876–891.

Falk JH (2016) *Identity and the museum visitor experience* (Walnut Creek, CA: Left Coast Press).

Falk JH, Koran Jr JJ, Dierking LD, Dreblow L (1985) Predicting visitor behavior. *Curator: The Museum Journal* 28(4):249–258.

Gallego G, Li A, Truong VA, Wang X (2020) Approximation algorithms for product framing and pricing. *Operations Research* 68(1):134–160.

Gallino S, Moreno A (2019) *Operations in an omnichannel world* (Springer).

Ghose A, Ipeirotis PG, Li B (2014) Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science* 60(7):1632–1654.

Ghose A, Li B, Liu S (2019) Mobile targeting using customer trajectory patterns. *Management Science* 65(11):5027–5049.

Giannini T, Bowen JP (2019) Museums and digital culture: New perspectives and research .

Greenberger A (2021) The louvre's looted renaissance masterpiece: New book explores the plundering of a veronese painting. *ARTnews* May 10:online, URL `https://www.artnews.com/art-news/news/veronese-wedding-feast-at-cana-plunder-book-review-1234592381/`.

Güler ÖK (2009) *A simulation application for visitor circulation in exhibition environments*. Ph.D. thesis, Bilkent University.

Guo X, Grushka-Cockayne Y, De Reyck B (2021) Forecasting airport transfer passenger flow using real-time data and machine learning. *Manufacturing & Service Operations Management* .

He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Hui SK, Bradlow ET, Fader PS (2009a) Testing behavioral hypotheses using an integrated model of grocery store shopping path and purchase behavior. *Journal of consumer research* 36(3):478–493.

Hui SK, Fader PS, Bradlow ET (2009b) Path data in marketing: An integrative framework and prospectus for model building. *Marketing Science* 28(2):320–335.

Hui SK, Inman JJ, Huang Y, Suher J (2013) The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies. *Journal of Marketing* 77(2):1–16.

Kabra A, Belavina E, Girotra K (2020) Bike-share systems: Accessibility and availability. *Management Science* 66(9):3803–3824.

King MJ, O'boyle J (2011) The theme park: The art of time and space. *Disneyland and culture: Essays on the parks and their influence* 5–18.

Kostami V, Ward AR (2009) Managing service systems with an offline waiting option and customer abandonment. *Manufacturing & Service Operations Management* 11(4):644–656.

Lanir J, Kuflik T, Dim E, Wecker AJ, Stock O (2013) The influence of a location-aware mobile guide on museum visitors' behavior. *Interacting with Computers* 25(6):443–460.

Larson JS, Bradlow ET, Fader PS (2005) An exploratory look at supermarket shopping paths. *International Journal of research in Marketing* 22(4):395–414.

Li J, Xu L, Tang L, Wang S, Li L (2018) Big data in tourism research: A literature review. *Tourism Management* 68:301–323.

Lu Y, Musalem A, Olivares M, Schilkrut A (2013) Measuring the effect of queues on customer purchases. *Management Science* 59(8):1743–1763.

Martínez-de Albéniz V, Valdivia A (2019) Measuring and exploiting the impact of exhibition scheduling on museum attendance. *Manufacturing & Service Operations Management* 21(4):761–779.

Melton A (1972) Visitor behavior in museums: Some early research in environmental design. *Human Factors* 14(5):393–403.

Milgram S (1970) The experience of living in cities. *Science* .

Montgomery AL, Li S, Srinivasan K, Liechty JC (2004) Modeling online browsing and path analysis using clickstream data. *Marketing Science* 23(4):579–595.

Müller-Trede J, Choshen-Hillel S, Barneron M, Yaniv I (2018) The wisdom of crowds in matters of taste. *Management Science* 64(4):1779–1803.

Nielsen L (1946) A technique for studying the behavior of museum visitors. *Journal of Educational Psychology* 37(2):103.

Pancras J, Sriram S, Kumar V (2012) Empirical investigation of retail expansion and cannibalization in a dynamic environment. *Management Science* 58(11):2001–2018.

Pluchino A, Garofalo C, Inturri G, Rapisarda A, Ignaccolo M (2013) Agent-based simulation of pedestrian behaviour in closed spaces: a museum case study. *arXiv preprint arXiv:1302.7153* .

Rajaram K, Ahmadi R (2003) Flow management to optimize retail profits at theme parks. *Operations Research* 51(2):175–184.

Ratner RK, Kahn BE, Kahneman D (1999) Choosing less-preferred experiences for the sake of variety. *Journal of Consumer Research* 26(1):1–15.

Reber R, Schwarz N, Winkielman P (2004) Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review* 8(4):364–382.

Serrell B, Aquarium JGS (1977) Survey of visitor attitude and awareness at an aquarium. *Curator: The Museum Journal* 20(1):48–52.

Sorensen H (2003) The science of shopping. *Marketing Research* 15(3):30–30.

Stokols D (1976) The experience of crowding in primary and secondary environments. *Environment and behavior* 8(1):49–86.

Suri R, Monroe KB (2003) The effects of time constraints on consumers' judgments of prices and products. *Journal of Consumer Research* 30(1):92–104.

Swann Jr WB, Pelham B (2002) Who wants out when the going gets good? psychological investment and preference for self-verifying college roommates. *Self and Identity* 1(3):219–233.

Thaler RH (1999) Mental accounting matters. *Journal of Behavioral Decision Making* 12(3):183–206.

Thiesse F, Fleisch E (2008) On the value of location information to lot scheduling in complex manufacturing processes. *International Journal of Production Economics* 112(2):532–547.

Tröndle M, Greenwood S, Bitterli K, van den Berg K (2014) The effects of curatorial arrangements. *Museum Management and Curatorship* 29(2):140–173.

Ursu RM (2018) The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science* 37(4):530–552.

Van Gogh Museum (2020) Multimedia Guide. `https://www.vangoghmuseum.nl/en/visit/multimedia-guide`, URL `https://www.vangoghmuseum.nl/en/visit/multimedia-guide`.

Veeraraghavan S, Debo L (2009) Joining longer queues: Information externalities in queue choice. *Manufacturing & Service Operations Management* 11(4):543–562.

Véron E, Levasseur M (1989) *Ethnographie de l'exposition: l'espace, le corps et le sens* (Bibliothèque publique d'information du Centre Pompidou).

Wathieu L (1997) Habits and the anomalies in intertemporal choice. *Management Science* 43(11):1552–1563.

Weingarten E, Bhatia S, Mellers B (2019) Multiple goals as reference points: One failure makes everything else feel worse. *Management Science* 65(7):3337–3352.

West MI (2011) *Disneyland and culture: Essays on the parks and their influence* (McFarland).

Xu L, Duan JA, Whinston A (2014) Path to purchase: A mutually exciting point process model for online advertising and conversion. *Management Science* 60(6):1392–1412.

Yoshimura Y, Sobolevsky S, Ratti C, Girardin F, Carrascal JP, Blat J, Sinatra R (2014) An analysis of visitors' behavior in the louvre museum: A study using bluetooth data. *Environment and Planning B: Planning and Design* 41(6):1113–1131.

Appendix

A Data Samples and Descriptive Statistics

Variable	Value	Variable	Value
id	890c2b5697c700000000_1547726853535	stop_id	40
language	fr	artwork	Van Gogh's Palette
tour	leisure	event_start_time	2019-01-17 13:15:49
tour_id	1	event_end_time	2019-01-17 13:16:30
floor_number	0	segment_completed	1

Table 7: An observation in the visitor log dataset.

Language	Proportion of visitors
English	32.78%
Italian	13.53%
French	8.61%
German	8.28%
Dutch	7.98%
Spanish	7.94%
Chinese	7.70%
Russian	4.68%
Portuguese	3.86%
Korean	3.11%
Japanese	1.50%

Table 8: Visitors in the Van Gogh Museum

Variable	Value	Variable	Value
artwork	The Potato Eaters	media	oil, canvas
details	1885, Nuenen, Netherlands	location	Van Gogh Museum, Amsterdam, Netherlands
style	Realism	dimensions	82cm × 114cm
genre	genre painting	tags	countryside, mealtimes

Table 9: An observation in the Wikiart.org dataset.

B Visitor Movement.

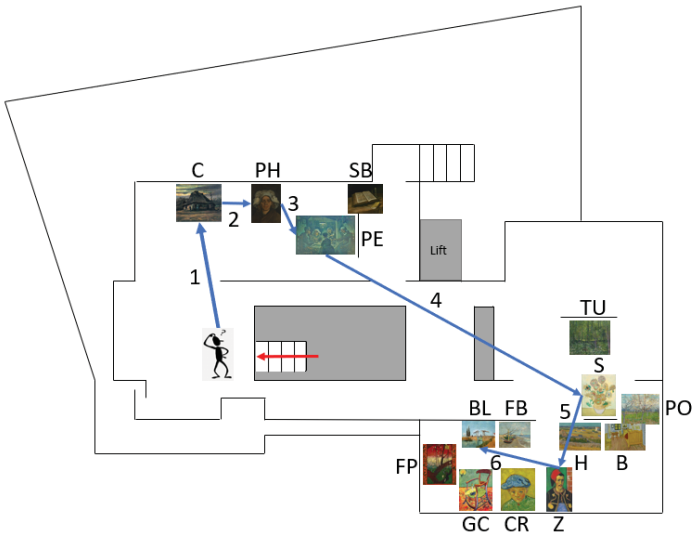

C: The Cottage; PH: Peasant Heads; SB: Still-life with Bible; PE: The Potato Eaters;
TU: Trees & Undergrowth; S: Sunflowers; PO: The Pink Orchard; H: The Harvest; B: The Bedroom
Z: The Zouave; PC: Portrait of Camille Roulin; GC: Gauguin's Chair; FP: Flowering Plum Orchard;
BL: Bridge of Langlois; FB: Fishing Boats on the Beach at Les Saint-Maries-de-la-Mer

Figure 7: Path taken by a randomly selected visitor on Floor 1 during October 2019.

Path segment	From Artwork	From Longitude	From Latitude	To Artwork	To Longitude	To Latitude	Distance (in meters)
1	Stairs	4.88113	52.35849	The Cottage	4.88127	52.35832	21.43
2	The Cottage	4.88127	52.35832	Peasant Heads	4.88124	52.35831	2.81
3	Peasant Heads	4.88124	52.35831	The Potato Eaters	4.88118	52.35833	4.82
4	The Potato Eaters	4.88117	52.35833	Sunflowers	4.88071	52.35846	35.22
5	Sunflowers	4.88071	52.35846	The Zouave	4.88073	52.35853	7.92
6	The Zouave	4.88072	52.35853	Bridge of Langlois	4.88082	52.35850	7.21

Table 10: Distance travelled by the randomly selected visitor.

C Deep Residual Learning

To track the artistic attributes of artworks in our analyses, we make use of the recent advancements in computer vision models based on *Convolutional Neural Nets* (CNN) in applications involving image/face recognition, object detection, image classification, etc. CNNs provide us a way to measure the the most salient visual features of each artwork and express them in a standardized numerical form (Banerjee et al. 2022). In simple terms, CNNs apply a local filter to small blocks of pixels; such local transformations are applied to all the blocks forming the image, and the output is a modified, lower-dimensional representation of the original image. Researchers have observed that such convolutional layers are often associated with interpretable transformations of the original image and extractions of features (e.g., edges, curves in the first layer, then more complex structures like squares and circles, etc.). See He et al. (2016) for more

details. The eventual goal of this process, is to reduce the dimensionality of the input while extracting the most important features, and thereby reduce the computational cost associated with downstream learning tasks such as image classification. This output, in the form of a low-dimensional vector, provides a numerical representation of the focal image. Specifically, we use a Residual Neural Network (ResNet) developed by He et al. (2016) that has been demonstrated to carry out classification tasks with very low error. It operates with 152 convolutional layers as opposed to the 8 layers of AlexNet and the 22 layers of GoogleNet. However, we use a smaller version of this neural network architecture, ResNet50, which was pre-trained on a corpus of over 1.28 million images from ImageNet.org. As a result, it is able to differentiate between a wide range of images irrespective of noise such as lighting, positioning, and contrast.

Accordingly, in the R environment, we first load a high-resolution picture of a painting. Using standard encoding, this image is then converted to a matrix of 896×896 pixels of Red, Green, and Blue (RGB) values. By calling the ResNet50 model with pre-trained weights, we carry out a sequential reduction of the features in layers. As the final output, we get a 1000-dimensional vectorised representation of the artwork.

Upon getting a vectorised representation of each artwork, we estimate the cosine dissimilarity between these vectors to get the difference between the artistic attributes of pairs of artworks. Accordingly, for two artworks i and j represented by V_i and V_j in vector form, respectively, we specify the $ArtDistance_{ij}$ variable measuring the difference between their artistic attributes as follows:

$$ArtDistance_{ij} = 1 - \frac{< V_i \ , \ V_j >}{||V_i|| \times ||V_j||}.$$

D Additional Materials from Section 5

D.1 Engageability and clickability of artworks

Artwork	$Artwork_i$ estimate "Engageability"	$Artwork_j$ estimate "Clickability"
Almond Blossom	-1.83***	1.50***
Bridge of Langlois	-1.44***	0.16***
Copies after Millet	-2.86***	-0.07***
Fishing Boats on the beach at Les Saint-Maries-de-la-Mer	-6.30***	-4.77***
Flowering Plum Orchard	-5.16***	-3.12***
Gauguin: Van Gogh painting Sunflowers	-1.05***	1.02***
Gauguin's Chair	-1.49***	0.21***
Irises	-2.08***	0.16***
Peasant Heads	-3.65***	0.44***
Portrait of Camille Roulin	NA	-4.59***
Seascape near Les Saint-Marie-de-la-Mer	-1.00***	1.08***
Self-Portrait with Grey Felt Hat	2.01	-3.89***
Self-Portrait with Straw Hat	-3.97***	1.37***
Self-Portrait with Straw Hat and Pipe	NA	-3.96***
Still life with Bible	4.58***	-3.86***
Sunflowers	-2.14***	0.52***
The Bedroom	-2.58***	0.63***
The Cottage	-3.22***	2.84***
The Harvest	-23.24	-4.24***
The Pink Orchard	-1.85***	1.73***
The Potato Eaters	-2.53***	1.58***
The Yellow House	-1.58***	1.15***
The Zouave	-2.04***	0.49***
Theo Van Gogh's Collector's Cabinet	-0.54***	1.67***
Tree Roots	-3.73***	-0.17***
Van Gogh's Palette	-2.59***	1.96***
Wheatfield with a Reaper	-2.27***	-0.34***
Wheatfield with Crows	-3.60***	0.00

$^{*}p < 0.05;^{**}\ p < 0.01;^{***}\ p < 0.001$.

Table 11: Estimates of the artworks fixed effects and their relation to the constructs of engageability and clickability from Besbes et al. (2016).

D.2 Additional models

Table 12 reports variations of the models included in Table 1, with the same covariates. In Table 14 we replicate the main models using the duration that a visitor has already spent, instead of the number of artworks seen. Finally, Table 15 provides alternative formulations with respect to the distance metrics.

	$Transition_{ij}^{v}$							
	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Physical Layout Factors								
$IsNearestPhysicalDistance_{\ell_i\ell_j}^{v}$	3.397***	1.345***	0.563***	0.583***	0.351***	0.583***	0.603***	0.395***
	(0.006)	(0.007)	(0.008)	(0.008)	(0.004)	(0.005)	(0.005)	(0.004)
$PhysicalDistance_{\ell_i\ell_j}^{v}$		-0.041***	-0.293***	-0.287***	-0.300***	-0.278***	-0.282***	-0.296***
		(0.001)	(0.002)	(0.002)	(0.000)	(0.000)	(0.000)	(0.000)
$RoomChange_{\ell_i\ell_j}^{v}$		-0.451***	-2.158***	-2.318***	-2.407***	-2.418***	-2.400***	-2.386***
		(0.017)	(0.029)	(0.030)	(0.006)	(0.006)	(0.006)	(0.006)
$FloorChange_{\ell_i\ell_j}^{v}$		-1.937***	-6.321***	-6.836***	-6.112***	-6.745***	-7.028***	-6.123***
		(0.022)	(0.033)	(0.035)	(0.005)	(0.005)	(0.005)	(0.005)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times RoomChange_{\ell_i\ell_j}^{v}$			0.237***	0.236***	0.252***	0.232***	0.232***	0.248***
			(0.002)	(0.002)	(0.000)	(0.000)	(0.000)	(0.000)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times FloorChange_{\ell_i\ell_j}^{v}$			0.311***	0.314***	0.311***	0.304***	0.310***	0.308***
			(0.002)	(0.002)	(0.000)	(0.000)	(0.000)	(0.000)
$ReturnedToFloor_{\ell_i\ell_j}^{v}$			-0.993***	-0.855***	-0.762***	-0.720***	-0.653***	-0.727***
			(0.014)	(0.014)	(0.011)	(0.011)	(0.012)	(0.011)
Digital Layout Factors								
$TourChange_{\ell_j}^{v}$				-3.093***	-2.072***	-2.874***	-2.993***	-2.104***
				(0.018)	(0.013)	(0.015)	(0.016)	(0.013)
$MMTguideDistance_{\ell_i\ell_j}^{v}$				-0.012***	-0.015***	-0.013***	-0.014***	-0.014***
				(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Artistic Factors								
$ArtYearDistance_{ij}^{v}$					-0.012***	0.013***	0.069***	-0.013***
					(0.002)	(0.002)	(0.002)	(0.002)
$ArtSizeDistance_{ij}^{v}$					0.411***	0.566***	0.639***	0.416***
					(0.005)	(0.005)	(0.005)	(0.005)
$ArtTagsDistance_{ij}^{v}$					-0.074***	-0.037***	-0.036***	-0.067***
					(0.003)	(0.003)	(0.003)	(0.003)
Environmental Factors								
$ArtSeen^{v}$						-0.117***	-0.206***	
						(0.001)	(0.001)	
$ArtSeen\_Q1^{v}$								6.080***
								(0.125)
$ArtSeen\_Q2^{v}$								4.401***
								(0.034)
$ArtSeen\_Q3^{v}$								2.903***
								(0.024)
$ArtSeen\_Q4^{v}$								2.624***
								(0.010)
$ArtSeen^{v} \times IsHighlight_{j}^{v}$							0.115***	
							(0.001)	
$LogCongestion^{v}$						0.358***	0.565***	
						(0.014)	(0.014)	
$LogCongestion\_Q1^{v}$								-0.042*
								(0.018)
$LogCongestion\_Q2^{v}$								0.051**
								(0.017)
$LogCongestion\_Q3^{v}$								0.032
								(0.017)
$LogCongestion\_Q4^{v}$								0.000
								(0.017)
$LogCongestion^{v} \times IsHighlight_{j}^{v}$							-0.240***	
							(0.008)	
$Thoroughness^{v}$						3.096***	3.096***	3.112***
						(0.011)	(0.011)	(0.010)
Fixed Effects								
$Artwork_{j}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$Artwork_{j}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$Language^{v}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$Hour^{v}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
$DayOfTheWeek^{v}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
AIC	176.27	182.53	188.57	192.61	198.60	204.64	210.66	216.61
Log-Likelihood	-1.72	-1.46	-1.42	-1.38	-1.39	-1.35	-1.33	-1.38
Num. parameters	89	92	95	97	100	103	106	109
Num. events	392,209	392,209	392,209	392,209	392,209	392,209	392,209	392,209
Num. obs.	6,221,248	6,221,248	6,221,248	6,221,248	6,221,248	6,221,248	6,221,248	6,221,248

*$p < 0.05$; $**p < 0.01$; $***p < 0.001$. Standard errors in parentheses.

Table 12: Pathway MNL estimation results (extended).

	$Transition_{ij}^{v}$	
	(13)	(14)
Physical Layout Factors		
$IsNearestPhysicalDistance_{\ell_i\ell_j}^{v}$	0.705***	0.694***
	(0.005)	(0.005)
$PhysicalDistance_{\ell_i\ell_j}^{v}$	-0.283***	-0.281***
	(0.000)	(0.000)
$RoomChange_{\ell_i\ell_j}^{v}$	-2.662***	-2.648***
	(0.006)	(0.006)
$FloorChange_{\ell_i\ell_j}^{v}$	-7.252***	-7.056***
	(0.005)	(0.005)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times RoomChange_{\ell_i\ell_j}^{v}$	0.241***	0.241***
	(0.000)	(0.000)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times FloorChange_{\ell_i\ell_j}^{v}$	0.315***	0.312***
	(0.000)	(0.000)
$ReturnedToFloor_{\ell_i\ell_j}^{v}$	-0.612***	-0.752***
	(0.012)	(0.012)
Digital Layout Factors		
$TourChange_{\ell_j}^{v}$	-2.950***	-3.081***
	(0.017)	(0.017)
$MMTguideDistance_{\ell_i\ell_j}^{v}$	-0.015***	-0.015***
	(0.000)	(0.000)
Artistic Factors		
$ArtYearDistance_{ij}^{v}$	0.061***	0.009***
	(0.002)	(0.002)
$ArtSizeDistance_{ij}^{v}$	0.531***	0.490***
	(0.005)	(0.005)
$ArtDistance_{ij}^{v}$	-1.606***	-1.602***
	(0.011)	(0.011)
Environmental Factors		
$ArtSeen^{v}$	-0.205***	
	(0.001)	
$ArtSeen\_Q1^{v}$		181.653***
		(1.902)
$ArtSeen\_Q2^{v}$		174.557***
		(0.033)
$ArtSeen\_Q3^{v}$		173.304***
		(0.024)
$ArtSeen\_Q4^{v}$		173.118***
		(0.010)
$ArtSeen^{v} \times IsHighlight_{j}^{v}$	0.111***	
	(0.001)	
$LogCongestion^{v}$	0.590***	
	(0.014)	
$LogCongestion\_Q1^{v}$		-0.049**
		(0.019)
$LogCongestion\_Q2^{v}$		0.037*
		(0.017)
$LogCongestion\_Q3^{v}$		0.027
		(0.017)
$LogCongestion\_Q4^{v}$		0.000
		(0.018)
$LogCongestion^{v} \times IsHighlight_{j}^{v}$	-0.244***	
	(0.008)	
$Thoroughness^{v}$	3.202***	3.012***
	(0.011)	(0.010)
Fixed Effects		
$Artwork_{j}$	Yes	Yes
$Artwork_{j}$	Yes	Yes
$Language^{v}$	Yes	Yes
$Hour^{v}$	Yes	Yes
$DayOfTheWeek^{v}$	Yes	Yes
AIC	210.65	216.67
Log-Likelihood	-1.34	-1.32
Num. parameters	106	109
Num. events	392,209	392,209
Num. obs.	6,221,248	6,221,248

$^{*}p < 0.05;^{**}p < 0.01;^{***}p < 0.001$. Standard errors in parentheses.

Table 13: Pathway MNL results with $ArtDistance_{ij}^{v}$ estimated by using the deep learning method of Appendix C.

	$Transition_{ij}^{v}$	
	(15)	(16)
Physical Layout Factors		
$IsNearestPhysicalDistance_{\ell_i\ell_j}^{v}$	0.593***	0.594***
	(0.005)	(0.005)
$PhysicalDistance_{\ell_i\ell_j}^{v}$	-0.283***	-0.281***
	(0.000)	(0.000)
$RoomChange_{\ell_i\ell_j}^{v}$	-2.464***	-2.411***
	(0.006)	(0.006)
$FloorChange_{\ell_i\ell_j}^{v}$	-6.991***	-6.905***
	(0.005)	(0.005)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times RoomChange_{\ell_i\ell_j}^{v}$	0.236***	0.233***
	(0.000)	(0.000)
$PhysicalDistance_{\ell_i\ell_j}^{v} \times FloorChange_{\ell_i\ell_j}^{v}$	0.311***	0.309***
	(0.000)	(0.000)
$ReturnedToFloor_{\ell_i\ell_j}^{v}$	-0.708***	-0.761***
	(0.012)	(0.012)
Digital Layout Factors		
$TourChange_{\ell_j}^{v}$	-3.105***	-3.179***
	(0.016)	(0.017)
$MMTguideDistance_{\ell_i\ell_j}^{v}$	-0.013***	-0.013***
	(0.000)	(0.000)
Artistic Factors		
$ArtYearDistance_{ij}^{v}$	0.040***	0.016***
	(0.002)	(0.002)
$ArtSizeDistance_{ij}^{v}$	0.624***	0.603***
	(0.005)	(0.005)
$ArtTagsDistance_{ij}^{v}$	-0.022***	-0.023***
	(0.003)	(0.003)
Environmental Factors		
$DurationUpto^{v}$	-0.867***	
	(0.005)	
$DurationUpto\_Q1^{v}$		3.049***
		(0.058)
$DurationUpto\_Q2^{v}$		1.215***
		(0.037)
$DurationUpto\_Q3^{v}$		0.492***
		(0.022)
$DurationUpto\_Q4^{v}$		0.000
		(0.010)
$DurationUpto^{v} \times IsHighlight_{j}^{v}$	0.080***	
	(0.002)	
$LogCongestion^{v}$	0.648***	
	(0.014)	
$LogCongestion\_Q1^{v}$		-0.105***
		(0.019)
$LogCongestion\_Q2^{v}$		0.015
		(0.017)
$LogCongestion\_Q3^{v}$		0.012
		(0.018)
$LogCongestion\_Q4^{v}$		0.000
		(0.018)
$LogCongestion^{v} \times IsHighlight_{j}^{v}$	-0.199***	
	(0.008)	
$Thoroughness^{v}$	3.480***	3.246***
	(0.011)	(0.010)
Fixed Effects		
$Artwork_{j}$	Yes	Yes
$Artwork_{j}$	Yes	Yes
$Language^{v}$	Yes	Yes
$Hour^{v}$	Yes	Yes
$DayOfTheWeek^{v}$	Yes	Yes
AIC	210.65	216.65
Log-Likelihood	-1.34	-1.34
Num. parameters	106	109
Num. events	392,209	392,209
Num. obs.	6,221,248	6,221,248

$^{*}p < 0.05$;** $p < 0.01$;*** $p < 0.001$. Standard errors in parentheses.

Table 14: Pathway MNL results with $DurationUpto^{v}$.

	$Transition_{ij}^{v}$	
	(17)	(18)
$IsNearestPhysicalDistance_{\ell_i\ell_j}^{v}$	2.692***	1.181***
	(0.003)	(0.007)
$PhysicalDistance_{\ell_i\ell_j}^{v}$		-0.073***
		(0.001)
$RoomChange_{\ell_i\ell_j}^{v}$		1.962***
		(0.033)
$FloorChange_{\ell_i=E,\ell_j=1}^{v}$		-3.358***
		(0.020)
$FloorChange_{\ell_i=E,\ell_j=2}^{v}$		-3.236***
		(0.037)
$FloorChange_{\ell_i=E,\ell_j=3}^{v}$		-3.492***
		(0.032)
$FloorChange_{\ell_i=0,\ell_j=0}^{v}$		1.206**
		(0.434)
$FloorChange_{\ell_i=0,\ell_j=1}^{v}$		0.262
		(0.434)
$FloorChange_{\ell_i=0,\ell_j=2}^{v}$		-0.983*
		(0.437)
$ReturnedToFloor_{\ell_i\ell_j}^{v}$		0.083***
		(0.017)
$ArtSeen^{v}$		-0.198***
		(0.003)
$LogCongestion^{v}$		0.518***
		(0.042)
Fixed Effects		
$Artwork_j$	No	Yes
$Artwork_j$	No	Yes
$Language^{v}$	No	Yes
$Hour^{v}$	No	Yes
$DayOfTheWeek^{v}$	No	Yes
AIC	-0.14	252.73
Log Likelihood	-2.14	-1.26
Num. parameters	1	127
Num. events	392,209	392,209
Num. obs.	6,221,248	6,221,248

$^{*}p < 0.05;^{**}p < 0.01;^{***}p < 0.001$. Standard errors in parentheses.

Interactions and factors related to MMT guide, artwork, and environment are omitted from this table for brevity.

$FloorChange_{i\geq 1,j=n}^{v}$, interactions and factors related to MMT guide, artwork, and environment are omitted from this table for brevity.

Table 15: Pathway MNL results for alternative models used in Section 5.2.

D.3 Predictive power

Model	Reference	Log-likelihood	Classification accuracy	Pair-level MAPE
Without layout factors	Table 1 Model 0	-2.31	23.45%	8.05%
With the nearest artwork	Table 1 Model 1	-1.72	48.68%	5.43%
Main model	Table 1 Model 3	-1.34	63.00%	3.95%
With floor to floor FE	Appx D.2 Table 15 Model 18	-1.27	64.06%	3.68%
With pair-level FE	Table 1 Model X	-1.08	67.20%	2.51%

Table 16: Out-of-sample predictive performance of P-MNL models. Note: Log-likelihood is normalized, i.e., we report the ratio of log-likelihood with the number of observations, and hence this figure reports the average value of the log-probability of the actual choice in the model.

D.4 The value of a sequential model

LOGIT model. What alternative method can be employed to directly predict $\Pr[i \in \sigma^{v,\pi}]$? Perhaps, the most natural alternative to predict $\Pr[i \in \sigma^{v,\pi}]$ is to view it as a binary classification task, with respect to the aggregate outcome "whether or not item i is hit by the visitor along her path". In this context, the class of logistic regression models, which we refer to as LOGIT, constitutes a relevant benchmark to estimate the hit rates.

	Hit_j^v		
	The Pink Orchard (1)	Sunflowers (2)	Irises (3)
(Intercept)	-1.999***	-2.082***	-2.802***
	(0.036)	(0.035)	(0.036)
$Congestion_j^v$	0.230**	0.072**	0.002
	(0.030)	(0.028)	(0.030)
Fixed Effects			
$Language^v$	Yes	Yes	Yes
$DayOfTheWeek^v$	Yes	Yes	Yes
AIC	36.78	36.70	36.44
Normalized LLH	-0.39	-0.35	-0.22
Num. obs.	145,974	174,078	309,442

$^{*}p < 0.1;^{**}p < 0.05;^{***}p < 0.01$. Standard errors in parentheses.

Table 17: Artwork level logistic models. For the sake of brevity, we present the results for only 3 artworks.

We note that the LOGIT benchmark offers more flexibility than P-MNL to predict the hit rates because distinct model parameters can be easily fitted for each item $i \in \mathcal{N}$ to capture a direct relationship between v and $\Pr[i \in \sigma^{v,\pi}]$. By contrast, this relationship is only indirectly captured by P-MNL through the distribution of sequential choices. Despite this advantage in terms of model expressiveness, LOGIT overlooks the sequential structure of each visit.

D.5 Natural experiments

For every pair (i, j), we denote P_{ijt} as the proportion of transitions from i to j on day t, out of all the visits to i. We then estimate the following auto-regressive specification $P_{ij,t} = \alpha_{ij} + \beta_{ijk} P_{ij,t-k} + \epsilon_{ijk}$ for $k = [7, 8]$. The model is trained with daily data from September and October 2019 and results are shown in Table 18. Hence, we subsequently use its estimated response $Control_{ijt}$ as our control for each pair of artworks (i, j) and day t during the post-intervention period.

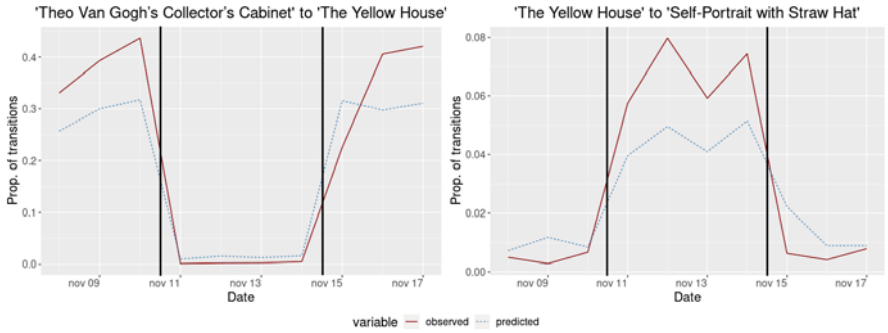

(a) For "The Yellow House".

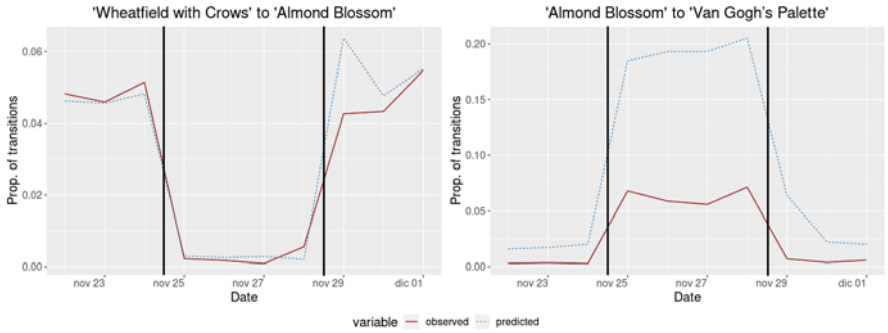

(b) For "Almond Blossom".

Figure 8: Predicted vs. observed transition probabilities during the relocation of a few artworks.

	$Observed_{ijt}$
(Intercept)	0.001***
	(0.0003)
$Observed_{ijt-7}$	0.514***
	(0.007)
$Observed_{ijt-8}$	0.451***
	(0.007)
R^2	0.91
Adj. R^2	0.91
Num. obs.	15,132

$^{*}p < 0.05;^{**}p < 0.01;^{***}p < 0.001$. Standard errors in parentheses.

Table 18: Estimating the control value. Subsample used is for transitions from or to the artwork "The Yellow House".

	$ObservedTreatment_{ijt}$				
	Nov. 11-14	Nov. 18-22	Nov. 25-28	Dec. 9-12	Overall
$PredictedTreatment_{ijt}$	0.927***	0.391***	0.639***	0.557***	0.741***
	(0.036)	(0.074)	(0.042)	(0.074)	(0.019)
Pair-level Control	Yes	Yes	Yes	Yes	Yes
R^2	0.571	0.458	0.559	0.427	0.267
Adj. R^2	0.485	0.258	0.408	0.305	0.220
Num. obs.	3527	1685	2137	3395	10744

$^{*}p < 0.1;^{**}\ p < 0.05;^{***}\ p < 0.01$.

Table 19: Relating the predicted and estimated treatment effects for four layout interventions across the full sample.

E Proof of Theorem 1

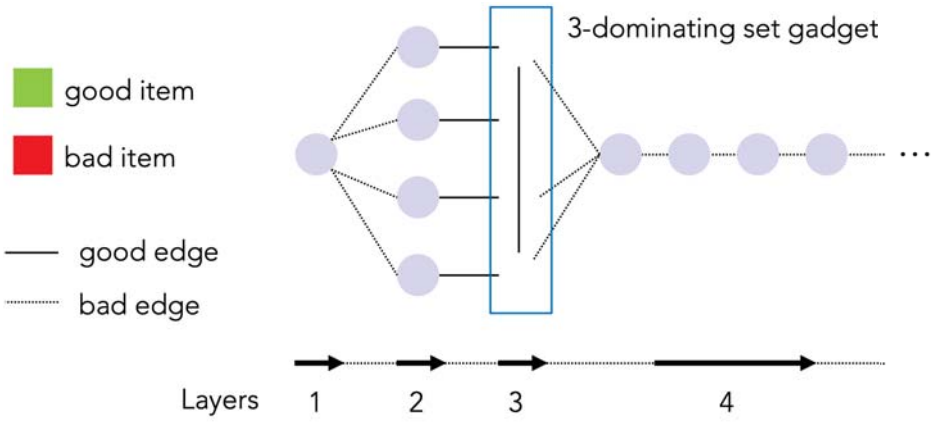

Figure 9: Pictorial representation of the hardness reduction.

We construct a reduction from the cubic dominating set problem, which is known to be APX-hard. Given a network (V, E) with $n = |V|$ vertices, the latter computational setting aims to distinguish whether there exists a dominating set of different cardinalities:

- *YES instance*: There exists a dominating set in (V, E) of cardinality k.
- *NO instance*: There exists no dominating set in (V, E) of cardinality $c \cdot k$ with $c > 1$.

Previous literature shows that the latter computational setting is NP-hard for some constant $c > 1$. Moreover, since the network is cubic, it follows that we can assume that there exists a family of NP-hard instances that further satisfy $k \geq \frac{n}{4}$; note that, when the latter condition is violated, the instance clearly falls under the "NO" case.

We proceed by devising a gap-preserving polynomial time reduction from the above decision problem to a computational task related to an approximation of the layout optimization problem. Specifically, we construct an instance of the layout optimization problem $\pi \mapsto \mathcal{D}_\pi$ by specifying a collection of items $\mathcal{N}$ with their corresponding nominal artwork-related utilities $u_a(\cdot, \cdot)$, and a network of locations $\mathcal{L}$ along with their corresponding edge utilities $u_l(\cdot, \cdot)$. Our family of instances is parameterized by an error parameter $\epsilon \in (0, 1)$.

- The collection $\mathcal{N} = G \cup B$ is formed by $5n$ good items $G = \{g_1, \ldots, g_{5n}\}$ and $n + 1$ bad items $B = \{b_1, \ldots, b_{n+1}\}$. Items are identical except for being good or bad.

- For all $i, j \in \mathcal{N}$, we decompose the artwork-related utility into two terms $u_a(i,j) = u_{1,1}(j) + u_{1,2}(i,j)$, where $u_{1,1}(j)$ is named the attractiveness and $u_{1,2}(i,j)$ is named the interaction. Good items $g \in G$ are more attractive than bad items: $u_{1,1}(g) = \log(\frac{1}{\epsilon}) \geq 0 = u_{1,1}(b)$. Moreover, good items $g, g' \in G$ have a strong positive interaction $u_{1,2}(g,g') = \log(\frac{n}{\epsilon})$, whereas bad items $b, b' \in B$ have a strong negative interaction $u_{1,2}(b,b') = -\log(\frac{n}{\epsilon})$. Finally, all mixed interactions are zero-valued, i.e., $u_{1,2}(g,b) = u_{1,2}(b,g) = 0$.

- The collection of locations $\mathcal{L}$ can be viewed as a graph with four layers: (1) A *distributor* star graph with a center ℓ^d connected to n distinct vertices $(\ell^d_v)_{v \in V}$; (2) A *copy* of V formed by the set of vertices $(\ell^c_v)_{v \in V}$; (3) A *path* $\ell^p_1, \ldots, \ell^p_q$ with $q = 4n$; (4) An *outside* vertex ℓ_0. These layers are connected amongst each other as follows. The center of the star graph ℓ^d is connected to the outside vertex ℓ_0. Each leaf of the distributor ℓ^d_v is connected to the vertex with similar index in the copy layer and all its neighbors, i.e., $\ell^c_v \cup \{\ell^c_{v'} : \{v', v\} \in E\}$. Now, each vertex ℓ^c_v of the copy layer is connected to ℓ^p_1 as well as the outside vertex ℓ_0. Finally, all vertices along the path are $\ell^p_1, \ldots, \ell^p_q$ with $q = 4n$ are connected to the outside vertex ℓ_0.

- Finally, it remains to specify the layout-related nominal utility function $u_l(\cdot, \cdot)$. In the first layer, the distributor utilities are zero-valued, i.e., $u_l(\ell^d, \ell^d_v) = 0$ for all $v \in V$, except with respect to the outside option $u_l(\ell^d, \ell_0) = \log(\frac{n^2}{\epsilon^3})$. In the second layer, any edge between a distributor leaf ℓ^d_v and a copy vertex $\ell^c_{v'}$ has a zero-valued utility $u_l(\ell^d_v, \ell^c_{v'}) = 0$. In the third layer, any edge between a copy vertex ℓ^c_v and ℓ^p_1 has a utility $u_l(\ell^c_v, \ell^p_1) = \frac{1}{2}\log(\epsilon)$, and the edge between ℓ^c_v and ℓ_0 has a large utility $u_l(\ell^c_v, \ell_0) = \log(\frac{1}{\epsilon})$. In addition, any edge between ℓ^c_v and another copy vertex $\ell^c_{v'}$ has a utility $u_l(\ell^c_v, \ell^c_{v'}) = \frac{1}{2}\log(\frac{1}{\epsilon})$. Finally, the utility of every edge (ℓ^p_k, ℓ^p_{k+1}) is zero-valued, i.e., $u_l = (\ell^p_k, \ell^p_{k+1}) = 0$, while $u_l(\ell^p_k, \ell_0) = \log(\frac{1}{\epsilon})$.

With these definitions at hand, let $\pi \mapsto \mathcal{D}_\pi$ be the mapping from item-to-location assignments to distributions over paths induced by the corresponding Pathway MNL instance. Our result proceeds from the following claims.

Claim 1. *Suppose that there exists an assignment π such that $\mathbb{E}_{\mathcal{D}_\pi}[|\sigma|] \geq \epsilon(4n-k) + 100\epsilon^{\frac{3}{2}}n + 100$. Then, there exists a dominating set in (V, E) of cardinality k.*

Proof. The proof proceeds by constructing a crucial upper bound on $\mathbb{E}_{\mathcal{D}_\pi}[|\sigma|]$. Let n_t be the number of good items in layer t for $t = 1, \ldots, 3$. Denote by $\tilde{n}_2$ the number of vertices in layer 2 (the copy) assigned or connected to a good item; in the dominating set terminology, the vertex in question is dominated.

Lemma 1. $\mathbb{E}[|\sigma^\pi|] \leq \frac{\epsilon \min\{n_1, \tilde{n}_2\}}{n} \cdot n_3 + 100\epsilon^{\frac{3}{2}}n + 100$.

Now, suppose that there exists no dominating set of cardinality k. To conclude the proof of Claim 1, it remains to upper bound the optimal solution of the following mathematical program:

$$\begin{aligned} \max_{n_1, n_2, \tilde{n}_2, n_3 \in [0,n]^4} \quad & \min\{n_1, \tilde{n}_2\} \cdot n_3 \\ \text{s.t.} \quad & n_1 + n_2 + n_3 \leq 5n \\ & n_2 + n - \tilde{n}_2 \geq k + 1 \ , \end{aligned}$$

where the objective corresponds to the right-hand side term in Lemma 1 and the first constraint follows from the total number of good items. To understand the last constraint, we observe that a dominating set of size $n_2 + n - \tilde{n}_2$ can be obtained by selecting all non-dominated vertices in layer 2 with respect to the good items assigned by π. By hypothesis, the cardinality of any dominating set is at least $k + 1$. Now, by noticing that $n_3 = 5n - n_1 - n_2$ at optimality, we can eliminate the constraints and formulate a relaxed mathematical program:

$$\max_{n_1, \tilde{n}_2 \in [0,n]^2} \quad \min\{n_1, \tilde{n}_2\} \cdot (6n - n_1 - \tilde{n}_2 - k - 1) \ .$$

Now, we can easily derive a continuous relaxation of this optimization setting:

$$\max_{\alpha \in [0,n]} \quad \alpha \cdot (6n - 2\alpha - k - 1) \ .$$

Using the fact that $c \cdot k \leq n$, we infer that the optimal value α^* for the variable α satisfies $\alpha = n$, which yields an upper bound of $n \cdot (4n - k)$. Combining this bound with Lemma 1, we obtain the desired contradiction:

$$\mathbb{E}[|\sigma^\pi|] < \epsilon(4n - k) + 100\epsilon^{\frac{3}{2}}n + 100 \ .$$

■

Proof of Lemma 1. Let $V^1 = \{v \in V : \pi^{-1}(\ell_v^d) \in G\}$ be the subset of vertex indices in layer 1 (i.e., the distributor) that hold a good item and let $V^2 = \{v \in V : \pi^{-1}(\ell_v^c) \in G \vee \mathcal{N}(\pi^{-1}(\ell_v^c)) \in G\}$ be the subset of vertex indices in layer 2 (i.e., the copy) that are dominated, meaning that either the location in question holds a good item or it has a neighbor in the copy that holds a good item. Let $\mathcal{L}^1 = \{\ell_v^d : v \in V^1\}$ and $\mathcal{L}^2 = \{\ell_v^d : v \in V^2\}$. Based on a straightforward path decomposition, we have

$$\begin{aligned}
&\mathbb{E}\left[|\sigma^\pi|\right] \\
&= \Pr\left[\pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] \cdot \mathbb{E}\left[|\sigma^\pi| \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] \\
&\quad + \Pr\left[\pi(\sigma_1^\pi) \notin \mathcal{L}^1\right] \cdot \mathbb{E}\left[|\sigma^\pi| \,|\, \pi(\sigma_1^\pi) \notin \mathcal{L}^1\right] \\
&\quad + \Pr\left[\pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \notin \mathcal{L}^2\right] \cdot \mathbb{E}\left[|\sigma^\pi| \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \notin \mathcal{L}^2\right] \\
&\leq \Pr\left[\pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] \cdot \mathbb{E}\left[|\sigma^\pi| \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] \\
&\quad + \frac{n}{n^2\epsilon^{-3} + n} \cdot 6n + \Pr\left[\pi(\sigma_1^\pi) \in \mathcal{L}^1\right] \cdot \mathbb{E}\left[|\sigma^\pi| \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \notin \mathcal{L}^2\right] \\
&\leq \Pr\left[\pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] \cdot \mathbb{E}\left[\sigma^\pi \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] \\
&\quad + 6\epsilon^3 + 24\epsilon^{\frac{3}{2}}n \ + 2\epsilon \\
&\leq \frac{\epsilon \min\{n_1, \tilde{n}_2\}}{n} \cdot \mathbb{E}\left[\sigma^\pi \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] \qquad (2) \\
&\quad + 6\epsilon^3 + 24\epsilon^{\frac{3}{2}}n + 2\epsilon \ ,
\end{aligned}$$

where the first inequality holds by noting that the user path is of length at most $6n$. Moreover, the transition probability $\Pr[\pi(\sigma_1^\pi) \notin \mathcal{L}^1]$ is maximized when $\mathcal{L}^1 = \emptyset$, meaning that all distributor vertices hold bad items: in such a case, we straightforwardly compute $\Pr[\pi(\sigma_1^\pi) \notin \mathcal{L}^1] = \frac{n}{n^2\epsilon^{-3}+n}$. The next inequality proceeds by noting that

$$\Pr\left[\pi(\sigma_1^\pi) \in \mathcal{L}^1\right] \leq \frac{n\epsilon^{-2}}{n\epsilon^{-2} + n\epsilon^{-3}} \leq \epsilon$$

where the first inequality follows from the observation that the probability on the left-hand side is maximized when $\mathcal{L}^1 = V$. In addition, we note that

$$\begin{aligned}
\mathbb{E}\left[|\sigma^\pi| \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \notin \mathcal{L}^2\right] &\leq 2 + \Pr\left[\sigma_3^\pi \neq 0 \,|\, \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \notin \mathcal{L}^2\right] \cdot 6n \\
&\leq 2 + 24\epsilon^{1/2}n \ ,
\end{aligned}$$

where the first inequality proceeds by noting that, after stage 2, the user transitions either to the outside option or towards a different vertex. In the former case, the path length is exactly

2, while in the latter case the path length is at most $6n$. The second inequality holds since

$$\begin{aligned}\Pr\left[\sigma_3^\pi = 0 \middle| \pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \notin \mathcal{L}^2\right] &= \Pr\left[\sigma_3^\pi = 0 \middle| \pi(\sigma_2^\pi) \notin \mathcal{L}^2\right] \\ &\geq \frac{\epsilon^{-1}}{3\epsilon^{-\frac{1}{2}} + \epsilon^{1/2}\epsilon^{-1} + \epsilon^{-1}} \\ &\geq 1 - 4\epsilon^{1/2} \ ,\end{aligned}$$

where the first inequality proceeds from the facts that the conditional probability is minimized when the first path vertex holds a good item and $\pi(\sigma_2^\pi)$ is only connected to vertices in layer 2 holding bad items.

Finally, inequality (2) is a consequence of the fact that

$$\begin{aligned}\Pr\left[\pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\right] &= \sum_{v \in V^1 \cap V^2} \Pr\left[\pi(\sigma_1^\pi) = \ell_v^d\right] \\ &= |V^1 \cap V^2| \cdot \frac{n\epsilon^{-2}}{n^2\epsilon^{-3} + |V^1| \cdot n\epsilon^{-2} + n - |V^1|} \\ &\leq \frac{\epsilon \min\{n_1, \tilde{n}_2\}}{n} \ ,\end{aligned}$$

where the second equality holds since all vertices in $\mathcal{L}^1$ hold good items. The inequality proceeds by noting that $|V^1 \cap V^2| \leq \min\{n_1, \tilde{n}_2\}$.

The remainder of our proof aims to upper bound $\mathbb{E}[|\sigma^\pi||E]$ where $E = \{\pi(\sigma_1^\pi) \in \mathcal{L}^1 \wedge \pi(\sigma_2^\pi) \in \mathcal{L}^2\}$. Define τ as the stopping time corresponding to a transition to the first path vertex ℓ_1^p or to the outside option, whichever is reached first, i.e., $\tau = \min\{k \geq 0 : \sigma_k^\pi \in \{0, \pi^{-1}(\ell_1^p)\}\}$. We distinguish between two cases depending on whether a good or a bad item is assigned to the first path vertex.

Case 1: $\boldsymbol{\pi^{-1}(\ell_1^p) \in B}$. The important observation is that $\Pr[\sigma_\tau^\pi = \pi^{-1}(\ell_1^p)|E] \leq \epsilon^{\frac{3}{2}}$. To explain this inequality, we note that until stage τ, the ratio between a transition to location ℓ_1^p and a transition to ℓ_0 is upper bounded by $\epsilon^{\frac{3}{2}}$; this bound corresponds to the case where the current location holds a good item, and thus, the attractiveness of $\pi^{-1}(\ell_1^p)$ is exactly $\epsilon^{\frac{1}{2}}$. Moreover, conditional on reaching ℓ_1^p, the remaining path length is clearly upper bounded by $4n$. By combining these observations, we have

$$\begin{aligned}&\mathbb{E}\left[|\sigma^\pi| \middle| E\right] \\ &= \mathbb{E}\left[\tau \middle| E\right] + \Pr\left[\sigma_\tau^\pi = \pi^{-1}(\ell_1^p) \middle| E\right] \cdot \mathbb{E}\left[|\sigma^\pi| - \tau \middle| \sigma_\tau^\pi = \pi^{-1}(\ell_1^p)\right] \\ &\leq 2 + n + 4\epsilon^{\frac{3}{2}} n \\ &\leq 2 + n_3 + 4\epsilon^{\frac{3}{2}} n \ ,\end{aligned}$$

where the first inequality proceeds from the bounds described above. The last inequality holds given that there $5n$ good items and only $2n+1$ locations in the first two layers, implying that $n_3 \geq 3n - 1 \geq n$ for a sufficiently large value of n.

Case 2: $\boldsymbol{\pi^{-1}(\ell_1^p) \in G}$. Here, we have

$$\begin{aligned}&\mathbb{E}\left[|\sigma^\pi| \middle| E\right] \\ &= \mathbb{E}\left[\tau - 1 \middle| E\right] + \Pr\left[\sigma_\tau^\pi = \pi^{-1}(\ell_1^p) \middle| E\right] \cdot \left(1 + \sum_{k=1}^{4n} \prod_{q=1}^{k} \Pr\left[\sigma_{\tau+q}^\pi = \pi^1(\ell_{q+1}^p) \middle| \sigma_{\tau+q-1}^\pi = \pi^1(\ell_q^p)\right]\right) \\ &= \mathbb{E}\left[\tau \middle| E\right] + \Pr\left[\sigma_\tau^\pi = \pi^{-1}(\ell_1^p) \middle| E\right] \cdot \left(1 + \sum_{k=1}^{4n-1} \prod_{q=1}^{k} \frac{\epsilon w_q}{1 + \epsilon w_q}\right) - 1 \ ,\end{aligned}$$

where the first equality proceeds from the path structure and the fact that each vertex can be visited at most once; these properties entail that there are only forward transitions along the path $\ell_1^p, \ell_2^p, \ldots$ until reaching the outside option. In the next equation, we define the coefficients $(w_k)_{k\in[4n-1]}$ as follows:

$$w_k = \begin{cases} \frac{n}{\epsilon^2} & \text{if } \pi^{-1}(\ell_k^p), \pi^{-1}(\ell_{k+1}^p) \in G \ , \\ \frac{1}{\epsilon} & \text{if } \pi^{-1}(\ell_k^p) \in B \text{ and } \pi^{-1}(\ell_{k+1}^p) \in G \ , \\ 1 & \text{if } \pi^{-1}(\ell_k^p) \in G \text{ and } \pi^{-1}(\ell_{k+1}^p) \in B \ , \\ \frac{\epsilon}{n} & \text{if } \pi^{-1}(\ell_k^p), \pi^{-1}(\ell_{k+1}^p) \in B \ . \end{cases}$$

Let k^* be the $\lceil \log n \rceil$-th index $k \in [4n]$ with a small transition probability $\frac{\epsilon w_k}{1+\epsilon w_k} \leq \frac{1}{2}$. In the above case disjunction, this inequality is satisfied for all transitions except between consecutive good items. By definition, for every $k \geq k^*$, we have

$$\sum_{k=k^*}^{4n-1} \prod_{q=1}^{k} \frac{\epsilon w_q}{1+\epsilon w_q} \leq \sum_{k=k^*}^{4n-1} \frac{1}{n} \leq 4 \ .$$

Additionally, we note that $k^* \leq n_3 + \lceil \log n \rceil$ since the first k^* locations in the path are either filled by good items or by at most $\lceil \log n \rceil$ bad items. By combining the above inequalities, we infer that

$$\mathbb{E}\left[\, |\sigma^\pi| \,|\, E\right] \leq \mathbb{E}\left[\, \tau \,|\, E\right] + 5 + n_3 + \lceil \log n \rceil \leq n_3 + 7 + 4\epsilon^{-1} + \lceil \log n \rceil \ ,$$

where the last inequality proceeds from the next claim.

Claim 2. $\mathbb{E}[\tau|E] \leq 2 + 4\epsilon^{-1}$.

Proof. The desired claim proceeds by providing a lower bound on the conditional probability of transitioning to locations ℓ_0 or ℓ_1^p $k = 3, 4, \ldots, \tau$. From stage $k = 3$ onwards, until reaching stage t, we are located in layer 2 of the graph (the copy). Suppose that the current location holds a good item. The probability of transitioning to layer 2 is maximized if the current location is connected to three copy vertices assigned with good items. Based on this reasoning, we have

$$\Pr\left[\, \sigma_k^\pi \in \{0, \pi^{-1}(\ell_1^p)\} \,\middle|\, \sigma_{k-1}^\pi \in G, S_k \right] \geq \frac{n\epsilon^{-\frac{3}{2}} + \epsilon^{-1}}{3n\epsilon^{-\frac{5}{2}} + n\epsilon^{-\frac{3}{2}} + \epsilon^{-1}} \geq \frac{\epsilon}{4} \ . \tag{3}$$

Conversely, suppose that the current location holds a bad item. Similarly to the previous case, the probability of transitioning to layer 2 is maximized if the current location is connected to three copy vertices assigned with good items. Here, we have

$$\Pr\left[\, \sigma_k^\pi \in \{0, \pi^{-1}(\ell_1^p)\} \,\middle|\, \sigma_{k-1}^\pi \in G, S_k \right] \geq \frac{\epsilon^{-1} + \epsilon^{-\frac{1}{2}}}{3\epsilon^{-\frac{3}{2}} + \epsilon^{-1} + \epsilon^{-\frac{1}{2}}} \geq \frac{2}{5}\epsilon \geq \frac{\epsilon}{4} \ . \tag{4}$$

By combining inequalities (3)-(4), we infer that

$$\mathbb{E}\left[\, \tau \,|\, E\right] \leq 2 + \frac{\epsilon}{4} \cdot \sum_{k=0}^{n} \left(1 - \frac{\epsilon}{4}\right)^{k-1} k \leq 2 + 4\epsilon^{-1}$$

■

■

Claim 3. *Suppose that there exists no assignment π such that $\mathbb{E}_{\mathcal{D}_\pi}[|\sigma|] \geq \epsilon\,(4n - \lceil ck \rceil) - 120\epsilon^{\frac{3}{2}}n$. Then, there exists no dominating set in (V, E) of cardinality $\lceil ck \rceil$.*

Proof. Suppose ad absurdum that there exists a dominating set V^* in (V, E) of cardinality smaller or equal to $\lceil ck \rceil$. We define the assignment π as follows:

- *Items of G:* Good items are allocated to locations ℓ^d_c and $\{\ell^d_v : v \in V\}$ in the distributor, $\{\ell^c_v : v \in V^*\}$ in the copy, and $\{\ell^p_1, \ldots, \ell^p_{4n-1-\lceil ck \rceil}\}$
- *Items in B:* All remaining locations are filled with bad items.

The proof proceeds by lower bounding the expected path length. The probability of transitioning from location ℓ^d_c to any vertex in $\{\ell^d_v : v \in V\}$ is exactly

$$\Pr\left[\sigma^\pi_1 \neq 0\right] = \frac{n^2\epsilon^{-2}}{n^2\epsilon^{-3} + n^2\epsilon^{-2}} \geq \epsilon(1-\epsilon)\ . \tag{5}$$

Next, we define τ as the stopping time corresponding to the first stage $k \geq 3$ where $\sigma^\pi_\tau \in \{0, \pi^{-1}(\ell^p_1)\}$ if $\sigma^\pi_1 \neq 0$ and $\tau = +\infty$ otherwise. We introduce the following probabilistic events: $E_1 = \{\sigma^\pi_3 \in \{0\} \cup B\}$ and $E_2 = \{\exists\ k \in [3, \tau] : \sigma^\pi_{k-1} \in G \wedge \sigma^\pi_k \in \{0\} \cup B\}$. The important observation is that, by the union bound, we have

$$\Pr\left[\sigma^\pi_\tau = 0 \middle|\, \sigma^\pi_1 \neq 0\right] \leq \Pr\left[E_1\right] + \Pr\left[E_2\right] \leq 8\sqrt{\epsilon}\ , \tag{6}$$

where the last inequality follows from Claims 4 and 5 below.

The last piece of our analysis consists in showing that, conditional on $\{\sigma^\pi_\tau = 0\}$, the probability of reaching the last good item along the path is sufficiently large. Namely,

$$\begin{aligned}
&\Pr\left[\sigma^\pi_{\tau+4n-2-\lceil ck \rceil} = \pi^{-1}\left(\ell^p_{4n-2-\lceil ck \rceil}\right) \middle|\, \sigma^\pi_\tau = \pi^{-1}(\ell^p_1)\right] \\
&= \prod_{k=1}^{4n-2-\lceil ck \rceil} \Pr\left[\sigma^\pi_{\tau+k} = \pi^{-1}(\ell^p_{1+k}) \middle|\, \sigma^\pi_{\tau+k-1} = \pi^{-1}(\ell^p_k)\right] \\
&= \prod_{k=1}^{4n-2-\lceil ck \rceil} \left(\frac{\epsilon^{-2}n}{\epsilon^{-2}n + \epsilon^{-1}}\right) \\
&\geq \left(1 - \frac{\epsilon}{n}\right)^{4n} \\
&\geq 1 - 4\epsilon\ .
\end{aligned} \tag{7}$$

Consequently, by combining (5)-(7), we obtain

$$\begin{aligned}
&\mathbb{E}\left[\sigma^\pi\right] \\
&\geq (4n - \lceil ck \rceil) \cdot \Pr\left[\sigma^\pi_{\tau+4n-2-\lceil ck \rceil} = \pi^{-1}\left(\ell^p_{4n-2-\lceil ck \rceil}\right)\right] \\
&\geq (4n - \lceil ck \rceil) \cdot \Pr\left[\sigma^\pi_{\tau+4n-2-\lceil ck \rceil} = \pi^{-1}\left(\ell^p_{4n-2-\lceil ck \rceil}\right) \middle|\, \sigma^\pi_\tau = \pi^{-1}(\ell^p_1)\right] \cdot \Pr\left[\sigma^\pi_\tau = \pi^{-1}(\ell^p_1)\right] \\
&\geq \epsilon\,(4n - \lceil ck \rceil) - 120\epsilon^{\frac{3}{2}}n\ .
\end{aligned} \tag{8}$$

Claim 4. $\Pr[E_1] \leq 4\sqrt{\epsilon}$.

Proof. Note that $\Pr[E_1] \leq \max\{\Pr[\sigma^\pi_3 \in \{0\} \cup B | \sigma^\pi_2 \in G], \Pr[\sigma^\pi_3 \in \{0\} \cup B | \sigma^\pi_2 \in B]\}$. Conditional on $\{\sigma^\pi_2 \in G\}$, the transition probability toward $\{0\} \cup B$ is maximized in the case where the current location is connected to copy vertices holding bad items, meaning that

$$\Pr\left[\sigma^\pi_3 \in \{0\} \cup B \middle|\, \sigma^\pi_2 \in G\right] \leq \frac{\epsilon^{-1} + 3\epsilon^{-\frac{1}{2}}}{\epsilon^{-1} + 3\epsilon^{-\frac{1}{2}} + n\epsilon^{-\frac{3}{2}}} \leq \frac{4\sqrt{\epsilon}}{n}\ . \tag{9}$$

Similarly, conditional on $\{\sigma_2^\pi \in B\}$, the transition probability toward $\{0\} \cup B$ is maximized in the case where the remaining neighbors of the current location in layer 2 hold bad items. Here, we have:

$$\Pr\left[\sigma_3^\pi \in \{0\} \cup B \middle| \, \sigma_2^\pi \in B\right] \leq \frac{\epsilon^{-1} + 2\epsilon^{\frac{1}{2}} n^{-1}}{\epsilon^{-1} + \epsilon^{-\frac{3}{2}} + 2\epsilon^{\frac{1}{2}} n^{-1} + \epsilon^{-\frac{1}{2}}} \leq 3\sqrt{\epsilon} \ .$$

■

Claim 5. $\Pr[E_2] \leq 4\sqrt{\epsilon}$.

Proof. Fix $k \in [3, n]$. We have

$$\Pr\left[\sigma_{k-1}^\pi \in G \wedge \sigma_k^\pi \in \{0\} \cup B\right] \leq \Pr\left[\sigma_k^\pi \in \{0\} \cup B \middle| \, \sigma_{k-1}^\pi \in G\right] \leq \frac{4\sqrt{\epsilon}}{n} \ ,$$

where the last inequality proceeds from a reasoning identical that of inequality (9). Since $\Pr[\tau \leq n] = 1$, the union bound implies that

$$\Pr[E_2] \leq \sum_{k=3}^{n} \Pr\left[\sigma_{k-1}^\pi \in G \wedge \sigma_k^\pi \in \{0\} \cup B\right] \leq 4\sqrt{\epsilon} \ .$$

■

■