Journal of Mathematical Psychology 131 (2026) 103016

Contents lists available at ScienceDirect

Journal of Mathematical Psychology

journal homepage: www.elsevier.com/locate/jmp

Structural identifiability of Order-of-Processing models: When do different cognitive architectures generate identical response time distributions

Aaron Khodami

Department of General Psychology, University of Padua, 35131, Padua, Italy

ARTICLE INFO

Keywords:
Order-of-Processing diagrams
Structural identifiability
Moment-generating functions
Cognitive architecture
Response-time distributions
Visual search

ABSTRACT

The Order-of-Processing (OP) method, introduced by Goldstein and Fisher in 1991, provides a powerful framework for deriving response-time distributions from stochastic network models of cognitive architecture. While the forward problem—predicting response-time distributions from a hypothesised network structure and process parameters—is well understood, the inverse problem remains unexamined: given an observed response-time distribution, when can researchers uniquely recover the underlying rate parameters, and when can they not? This work studies a restricted path-based subclass of OP models and develops a theory of structural identifiability for that subclass using moment-generating functions, Laplace-transform arguments, and diagram symmetries. The work shows that, under stated regularity and no-cancellation conditions, the multiset of rate parameters is generically recoverable from the reduced rational transform, whereas assignment of those rates to specific process labels can fail whenever the OP diagram contains nontrivial permutation symmetries. The work formalises the symmetry group of an OP diagram, show how its action induces observationally equivalent parameter assignments, and distinguish this within-framework non-identifiability from the broader serial–parallel mimicry problem. Applications to visual-search architectures illustrate symmetric, asymmetric, and partially symmetric cases and clarify the inferential scope of OPM-based analyses. The results provide practical guidance on when additional constraints or experimentally induced asymmetries are needed before fitted OPM parameters can be interpreted psychologically.

1. Introduction

The human cognitive system processes information through architectures that may involve multiple stages, partial-order constraints, parallel operations, stopping rules, and capacity limitations. Understanding the temporal organisation of these processes has been a central concern of cognitive psychology since the nineteenth century (Donders, 1868; Sternberg, 1969, 1984). The methodological transition from Donders' subtraction method to Sternberg's additive-factors method marked not only a change in technique, but also a change in the kinds of questions researchers sought to answer (Donders, 1868; Sternberg, 1969, 1984). Rather than merely estimating the duration of covert processes, investigators increasingly aimed to infer the underlying architecture of cognition: which subprocesses exist, how they are ordered, whether they operate serially or in parallel, and under what conditions they can be empirically distinguished (Townsend, 1972, 1984, 1990; Townsend & Ashby, 1983). This broader literature also established a caution that remains central to the present paper: similar response-time regularities can arise from structurally different systems, so inference from response times to architecture requires explicit mathematical conditions rather than informal intuition alone (Townsend, 1972, 1990; Townsend & Ashby, 1983).

Goldstein and Fisher (1991) introduced the Order-of-Processing method, hereafter referred to as the OPM, as a general mathematical framework for deriving response-time distributions from stochastic network models of cognitive architecture. Building on the earlier use of PERT and critical-path networks in psychology (Schweickert, 1978, 1980, 1982), and situated within the broader program of stochastic mental networks (Schweickert & Townsend, 1989; Townsend & Schweickert, 1989), the OPM provided a flexible formalism for representing cognitive structures that are neither purely serial nor purely parallel (Goldstein & Fisher, 1991, 1992). Within this critical-path tradition, Schweickert's latent-network and slack-based results were directed not only at representability, but also at recovering order relations among subprocesses from response-time manipulations and thereby treating subprocess order itself as an inferential target rather than a fixed assumption (Schweickert, 1978, 1982, 1992). This was an important step because the architectural space under consideration had already expanded well beyond simple serial-stage models: researchers were

E-mail address: aaron.a.khodami@gmail.com.

https://doi.org/10.1016/j.jmp.2026.103016
Received 19 December 2025; Received in revised form 6 May 2026; Accepted 31 August 2026
Available online 12 September 2026

concerned not only with serial versus parallel processing, but also with self-terminating versus exhaustive search, limited versus unlimited capacity, and stochastic dependence among process durations (Goldstein & Fisher, 1991; Townsend, 1984, 1990; Townsend & Ashby, 1983). In this setting, the OPM offered a tractable way to derive theoretically expected response-time distributions for a broad class of discrete, feedforward processing structures (Goldstein & Fisher, 1991, 1992).

The elegance of the OPM lies in its reduction of stochastic network problems to linear algebra. Specifically, the central theorem relates events in the OP diagram to linear combinations of process durations through the projection matrix

$$P = A_\alpha \left(A'_\alpha A_\alpha\right)^{-1} A'_\alpha \tag{1.1}$$

This reduction allows the derivation of response-time distributions in closed or semi-closed form for important classes of process-duration distributions (Goldstein & Fisher, 1991, 1992). It also clarifies why OP diagrams are more tractable than direct work on stochastic PERT networks: events defined by a path through the diagram can be translated into linear constraints on the underlying process durations or work requirements (Goldstein & Fisher, 1991, 1992). In that sense, the OPM preserves the stochastic-network perspective while substantially enlarging the class of structures for which exact distributional predictions can be obtained.

However, the foundational OPM literature is primarily concerned with the forward problem: given a specified structure and assumptions about process durations, what response-time distribution follows Goldstein and Fisher (1991, 1992)? That question is indispensable, but it does not by itself resolve the inverse problem that is crucial for empirical use. If an observed response-time distribution is compatible with more than one parameterisation or more than one structurally distinct assignment of parameters within the OPM, then any apparently precise architectural conclusion may be illusory. This inverse problem should be distinguished from the older serial–parallel mimicry literature. Classical results on identifiability and mimicking showed that distinct model classes can sometimes generate equivalent probability distributions on completion times (Townsend, 1972, 1990; Townsend & Ashby, 1983). The present paper does not claim to solve that broader problem for all cognitive architectures. Rather, it asks a narrower but still essential question: within the OPM itself, when does the response-time distribution uniquely determine the underlying parameter assignment and when does it not?

Fig. 1 clarifies this distinction. Panel (A) illustrates the kind of problem addressed in the present work: within a fixed OPM representation, different assignments of rate parameters to labelled processes may yield the same response-time transform. Panel (B) illustrates a different problem from the classical literature, namely serial–parallel mimicry across architectural families. The two issues are related in a broad inferential sense, but they are not the same mathematical problem.

This distinction matters because the OPM sits inside a much wider tradition of systems identification based on selective influence and factorial reasoning. Sternberg's original framework relied on the idea that different experimental factors could selectively affect different processes, but later work showed that this issue is more subtle once architectures become more complex and stochastic dependence among process durations is allowed (Schweickert & Townsend, 1989; Townsend, 1984; Townsend & Schweickert, 1989). Subsequent developments extended selective-influence reasoning beyond simple serial models and beyond stochastic independence, and connected it more explicitly to the general problem of systems identification (Dzhafarov, 2003; Schweickert, 1985; Schweickert et al., 2000; Townsend & Liu, 2022). These developments are directly relevant here. They show, first, that the empirical interpretation of response-time patterns depends on precise assumptions about which factors influence which subprocesses, and second, that reaction time alone is often not sufficient to resolve all architectural ambiguities. Accuracy, factorial contrasts, and richer distributional diagnostics may provide leverage unavailable to mean response times or unconstrained latency summaries (Schweickert, 1985; Schweickert et al., 2000; Townsend & Liu, 2022). The present study therefore addresses one specific inferential problem — structural identifiability within the OPM from response-time distributions — without claiming that response time alone can generally settle all questions of cognitive architecture.

It is also important to state the structural scope of the present analysis. Goldstein and Fisher's formulation makes clear that, for the finite total-order subclass of OP diagrams, one and only one path through the OP diagram is followed on any particular trial (Goldstein & Fisher, 1991, 1992). The results developed in the present paper are intended for that path-based, discrete, acyclic setting. They are therefore not a theory of all possible information-processing systems. In particular, they do not directly apply to continuous-flow or cascade conceptions in which information can be transmitted incrementally across processing levels before upstream processing is fully complete (McClelland, 1979). That contrast is important because the broader literature has shown that additive-factor style conclusions can look very different once contingent processing is allowed to overlap dynamically (McClelland, 1979).

Against that background, this paper develops a theory of structural identifiability for a class of OPM models using moment-generating functions and symmetry analysis.[1] The central claim is not that all non-uniqueness in cognitive architecture reduces to symmetry, nor that OPM-based models exhaust the relevant architecture space. Rather, the claim is that, within the OPM class analysed here, permutation symmetry in the diagram provides a fundamental source of non-identifiability. When parallel branches or structurally exchangeable processes occupy indistinguishable roles, multiple parameter assignments can imply the same response-time distribution. Conversely, when the relevant symmetries are absent, parameter recovery may be unique under the stated assumptions. Framing the problem this way provides both a theoretical clarification and a practical one: before fitting an OPM to data and interpreting its parameters psychologically, the researcher should determine whether the proposed diagram is structurally identifiable or whether additional constraints or experimental manipulations are required to break the symmetry.

The organisation of this paper is as follows. First, the mathematical foundations of structural identifiability in the OPM are developed through moment-generating-function representations of response-time distributions. Next, the paper distinguishes between recovery of the multiset of rate parameters and recovery of their assignment to specific processes, formalises the role of diagram symmetries in that distinction, and explains how the symmetry group can be computed in practice. It then introduces equivalence classes of OP diagrams and illustrates symmetric, asymmetric, and partially symmetric cases using visual-search examples. Finally, it discusses the implications of these results for experimental design, model interpretation, and the proper inferential scope of OPM-based analyses.

2. Mathematical foundations: Structural identifiability via moment-generating functions

The foundational connection between OP diagrams and moment-generating functions first requires establishing the relationship among response time, network structure, and process durations. Goldstein and Fisher (1991, 1992) showed that, on a given trial, when a specific path α through an OP diagram is realised, the total response time can be written as the sum of the state durations along that path. Each state's duration is determined by the processes completed during

[1] The analysis is algebraic only in the restricted sense that it relies on rational transforms, pole structure, and finite permutation groups. It does not invoke algebraic geometry in the modern Grothendieckian sense.

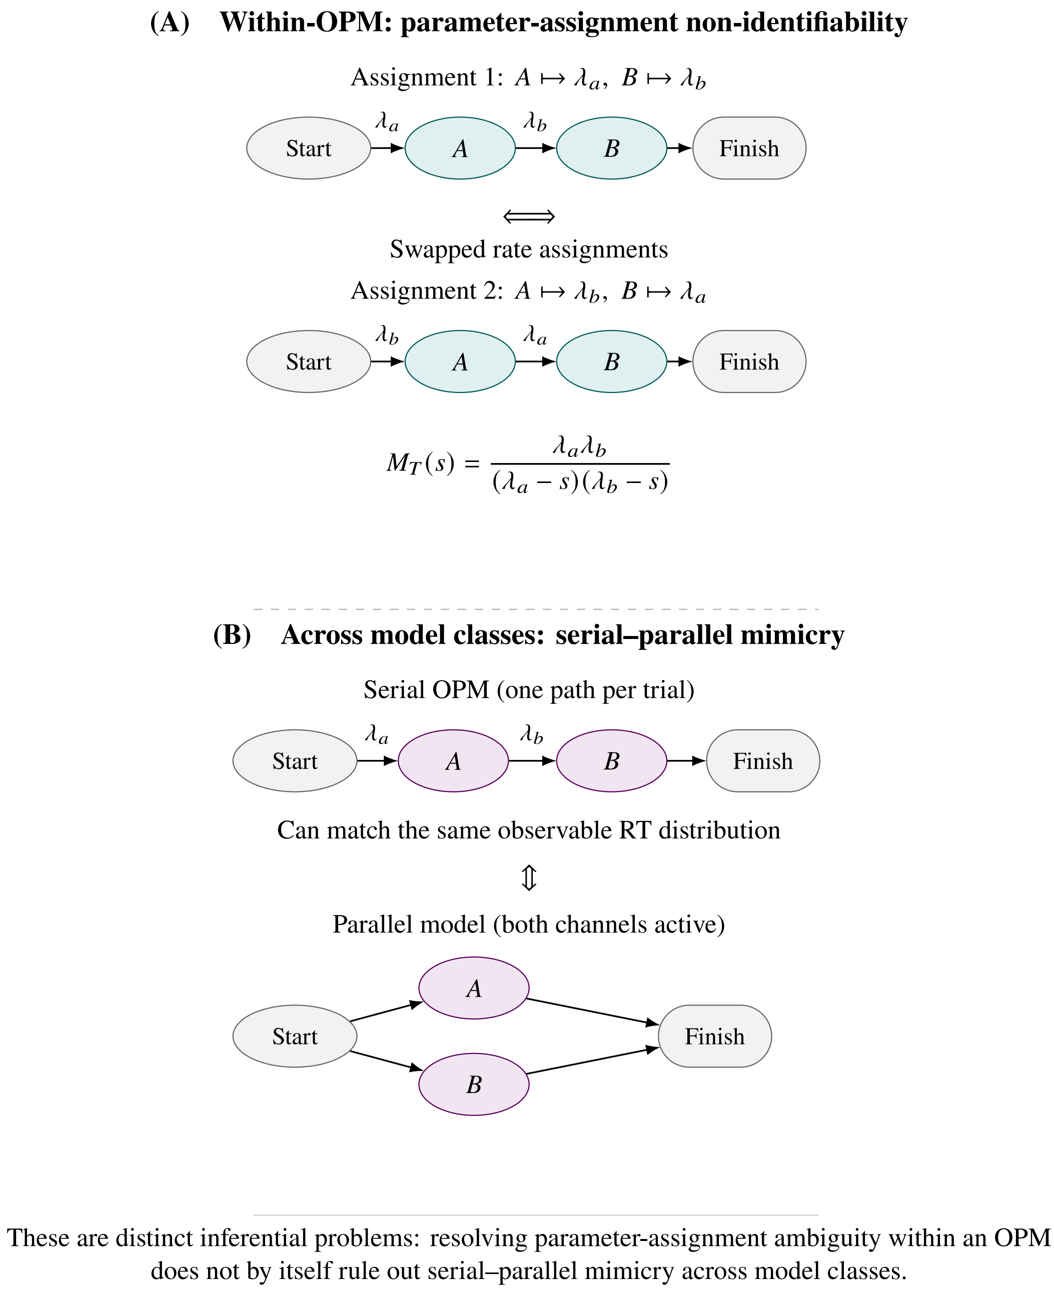

Fig. 1. Schematic distinction between two different sources of non-uniqueness in response-time modelling. (A) The upper panel illustrates parameter-assignment non-identifiability within a fixed Order-of-Processing model: different assignments of rate parameters to labelled processes can yield the same moment-generating function. (B) The lower panel illustrates the classical serial–parallel mimicry problem, in which models from different architectural families may generate the same observable response-time distribution (Townsend, 1972; Townsend & Ashby, 1983). The figure is intended as a conceptual clarification: resolving one problem does not by itself resolve the other.

the corresponding state transition. For exponential process durations, this pathwise structure yields a particularly tractable transform representation. However, the present development concerns a restricted subclass of OP models. Specifically, the work focuses on path-based OP models in which, on each trial, a realised path $W = \alpha$ determines the response time, and conditional on that realised path the response time is generated by a sum of state durations associated with that path. This restriction is narrower than the full OPM, but it is the subclass for which the transform-based identifiability analysis developed below can be stated clearly (Goldstein & Fisher, 1991, 1992).

Let $D = \langle \Pi, S, \Gamma \rangle$ be an Order-of-Processing diagram with exponential process durations $X_i \sim \mathrm{Exp}(\lambda_i)$, where $\lambda_i > 0$ is the rate parameter for process $x_i \in \Pi$. Let W denote the random realised path through the diagram, so that $W = \alpha$ means that path $\alpha \in \Gamma$ is followed on a given trial. Then, for values of s in a neighbourhood of 0 for which the transform exists, the moment-generating function of the response time T is

$$M_T(s) = \sum_{\alpha \in \Gamma} P(W = \alpha)\, M_{T|W=\alpha}(s), \tag{2.1}$$

where $P(W = \alpha)$ denotes the probability that the realised path is α, and $M_{T|W=\alpha}(s)$ is the conditional moment-generating function of response time given that α is followed. By the law of total expectation applied to the path variable W,

$$M_T(s) = E\left[e^{sT}\right] = \sum_{\alpha \in \Gamma} E\left[e^{sT} \mid W = \alpha\right] P(W = \alpha). \tag{2.2}$$

According to the OPM, when the realised path is α, the response time equals the sum of the state durations along that path,

$$T = \sum_{j=1}^{N(\alpha)} T_{\alpha_j}, \tag{2.3}$$

where T_{α_j} denotes the duration of state s_{α_j} and $N(\alpha)$ is the number of nonterminal states along path α (Goldstein & Fisher, 1991). By the same pathwise theorem, the state-duration vector satisfies

$$T_\alpha = \left(A'_\alpha A_\alpha\right)^{-1} A'_\alpha X_\alpha, \tag{2.4}$$

where A_α is the path-specific matrix induced by the OP-diagram structure and X_α is the vector of primitive process durations associated with path α. Thus, state durations are linear combinations of the underlying process durations. For a state duration of the form

$$T_{\alpha_j} = \sum_k c_{jk} X_k, \tag{2.5}$$

where the X_k are independent exponential variables, the corresponding moment-generating function is

$$M_{T_{\alpha_j}}(s) = \prod_k \frac{\lambda_k}{\lambda_k - c_{jk} s}, \tag{2.6}$$

provided all denominators remain non-zero.

For the restricted subclass considered here, we impose the additional simplifying assumption that, along any realised path α, the primitive process durations contributing to distinct states are disjoint. Under this condition, the state durations along α are independent, and the conditional response time reduces to a sum of independent exponential process durations:

$$T \mid W = \alpha = \sum_{i\in\alpha} X_i. \tag{2.7}$$

Hence

$$M_{T|W=\alpha}(s) = \prod_{i\in\alpha} \frac{\lambda_i}{\lambda_i - s}, \tag{2.8}$$

for all $s < \min_{i\in\alpha} \lambda_i$. Substituting Eq. (2.8) into Eq. (2.1) gives

$$M_T(s) = \sum_{\alpha\in\Gamma} P(W=\alpha) \prod_{i\in\alpha} \frac{\lambda_i}{\lambda_i - s}. \tag{2.9}$$

Eq. (2.9) is the basic transform representation used throughout the remainder of this section.

Proposition 2.1 (*Path-Mixture Transform*). *Under the path-based subclass described, with disjoint primitive process contributions across distinct states on each realised path, the response-time transform is given by Eq. (2.9).*

The path-mixture form in Eq. (2.9) implies that $M_T(s)$ is a rational function of s:

$$M_T(s) = \frac{P(s)}{Q(s)}, \tag{2.10}$$

for suitable polynomials $P(s)$ and $Q(s)$. Before any algebraic cancellation is carried out, one admissible common denominator is

$$Q(s) = \prod_{i=1}^{N} \left(\lambda_i - s\right)^{d_i}, \tag{2.11}$$

where $d_i \geq 0$ records how many path terms in Eq. (2.9) contain process i. Eq. (2.11) therefore gives a common denominator for the unreduced sum in Eq. (2.9). After numerator and denominator are reduced, some factors may cancel. Consequently, statements about poles and identifiability must be made with respect to the *reduced* rational form of $M_T(s)$, not the unreduced denominator. This representation is crucial because it establishes that the structure of the transform, and in particular the location of its poles in the reduced rational form, encodes information about the process rate parameters. The first identifiability question is therefore whether the transform uniquely determines the response-time distribution. A fundamental principle of probability theory states that, when the moment-generating function exists on an open interval containing 0, it uniquely determines the corresponding distribution (Curtiss, 1942; Stoyanov, 2015). Since response time is nonnegative, it is also useful to note that the Laplace transform

$$L_T(u) = E\left[e^{-uT}\right] = M_T(-u), \qquad u \geq 0, \tag{2.12}$$

uniquely determines the distribution of T under the present assumptions (Singhal & Vlach, 1975). Equivalently, the density function may be recovered by Laplace inversion:

$$f_T(t) = \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} e^{-st} M_T(s)\, ds, \tag{2.13}$$

where the contour is chosen in the usual way. Thus, if two models in the present subclass satisfy

$$M_T^{(1)}(s) = M_T^{(2)}(s) \tag{2.14}$$

for all s in a neighbourhood of 0, then they imply the same response-time distribution:

$$f_T^{(1)}(t) = f_T^{(2)}(t). \tag{2.15}$$

The next question is what aspect of the process parameters is recoverable from the transform. Under the path-mixture representation in Eq. (2.9), with distinct rate parameters and in the generic case where no pole cancellation occurs in the reduced rational form, the poles occur at the rate values λ_i that remain in the reduced denominator. Consequently, the *multiset* of rate parameters is generically structurally identifiable from the response-time distribution. This is an important result, but it is only a partial solution to the identifiability problem. Recovery of the multiset of poles does not by itself determine which rate belongs to which specific process label.

Proposition 2.2 (*Generic Multiset Identifiability*). *Under the path-mixture representation in Eq. (2.9), if the rate parameters are distinct and no pole cancellation occurs in the reduced rational form of $M_T(s)$, then the multiset of rate parameters appearing as poles in that reduced form is generically structurally identifiable from the response-time distribution.*

To address assignment of rates to processes, define the path-process incidence matrix I by

$$I_{k,i} = \begin{cases} 1 & \text{if process } x_i \text{ appears in path } \alpha_k, \\ 0 & \text{otherwise,} \end{cases} \tag{2.16}$$

for $k = 1,\dots,|\Gamma|$ and $i = 1,\dots,N$. Each row of I encodes which processes contribute to a particular realised path. Eq. (2.9) can then be written as

$$M_T(s \mid \lambda) = \sum_{k=1}^{|\Gamma|} P(W=\alpha_k) \prod_{i:\, I_{k,i}=1} \frac{\lambda_i}{\lambda_i - s}, \tag{2.17}$$

where

$$\lambda = \left(\lambda_1,\dots,\lambda_N\right) \tag{2.18}$$

denotes the vector of process-specific rate parameters. Now consider another assignment

$$\lambda' = \left(\lambda'_1,\dots,\lambda'_N\right). \tag{2.19}$$

Assignment non-identifiability occurs when distinct parameter vectors λ and λ' yield the same transform in Eq. (2.17). The source of this non-identifiability is most naturally expressed in terms of symmetry.

Definition 2.3 (*Symmetry Group*). Let $\sigma \in S_N$ be a permutation of process labels, acting on the parameter vector by

$$\sigma(\lambda) = \left(\lambda_{\sigma(1)},\dots,\lambda_{\sigma(N)}\right). \tag{2.20}$$

The symmetry group of the OP diagram is

$$\mathcal{G}(D) = \left\{\sigma \in S_N \,:\, M_T(s \mid \lambda) = M_T(s \mid \sigma(\lambda)) \text{ for all admissible } s\right\}. \tag{2.21}$$

Definition 2.3 captures those relabellings of process-specific rates that do not alter the response-time transform. In practice, a sufficient route to identifying elements of $\mathcal{G}(D)$ is to inspect the incidence matrix I together with the path probabilities $P(W = \alpha_k)$: if a relabelling preserves the structural role of the processes across all realised paths and leaves the path weights unchanged, then it belongs to $\mathcal{G}(D)$.

The computation is algorithmic. One may construct a coloured bipartite graph whose vertices represent paths and processes, whose edges encode the incidence relation in Eq. (2.16), and whose path vertices are coloured according to their path probabilities or probability-equivalence classes. The automorphism group of this graph yields the same symmetry information as $\mathcal{G}(D)$.

Theorem 2.4 (*Symmetry-Based Assignment Non-Identifiability*). *If the symmetry group $\mathcal{G}(D)$ contains a nonidentity permutation, then there exist distinct assignments of rate parameters that induce the same response-time transform and hence the same response-time distribution. Therefore, for generic parameter values,*

$$|\mathcal{G}(D)| > 1 \quad \Longrightarrow \quad \textit{assignment non-identifiability}. \tag{2.22}$$

This is the central symmetry-based source of non-identifiability in the present subclass of OPM models. Parameter vectors generating the same transform form orbits under the action of $\mathcal{G}(D)$, so the size of the symmetry group quantifies the severity of the ambiguity. A particularly important case is that of symmetric parallel OP diagrams. Suppose the path set contains K structurally identical and exchangeable paths $\alpha_1, \ldots, \alpha_K$ with equal path probabilities. Then

$$M_T(s) = \sum_{k=1}^{K} P(W = \alpha_k) \prod_{i \in \alpha_k} \frac{\lambda_i}{\lambda_i - s}, \tag{2.23}$$

and if $P(W = \alpha_k) = 1/K$ for all k, this becomes

$$M_T(s) = \frac{1}{K} \sum_{k=1}^{K} \prod_{i \in \alpha_k} \frac{\lambda_i}{\lambda_i - s}. \tag{2.24}$$

Any permutation that merely exchanges these structurally identical paths leaves the transform unchanged, because it preserves both the path probabilities and the collection of conditional path terms. Therefore,

$$M_T(s \mid \lambda) = M_T(s \mid \sigma(\lambda)) \tag{2.25}$$

for every permutation σ that respects the symmetry of the diagram. In such cases, the assignment of rates to specific path labels is not structurally identifiable.

Corollary 2.5 (*Trivial Symmetry Removes Permutation-Based Ambiguity*). *If $|\mathcal{G}(D)| = 1$, then permutation symmetry no longer provides a source of non-identifiability within the present path-mixture subclass.*

The positive content of Corollary 2.5 should not be overstated. A trivial symmetry group removes *permutation-based* ambiguity within the present path-mixture subclass; it does not by itself establish global uniqueness against every possible algebraic coincidence, nor does it settle the broader question of observational equivalence across distinct model classes. In particular, the present analysis concerns non-identifiability *within* a fixed subclass of OPM models. It is not equivalent to the classical problem of serial–parallel mimicry studied by Townsend (1972) and Townsend and Ashby (1983), where distinct architectural families may generate identical observable distributions. The distinction is essential: the present section identifies a structural source of ambiguity inside the model class under study, rather than claiming universal discriminability across all cognitive architectures.

This distinction also clarifies the role of Goldstein and Fisher's regularity conditions. In the original OPM, these conditions ensure that path probabilities are well defined, that no simultaneity disrupts the state structure, and that the response-time distribution can be derived from the diagram (Goldstein & Fisher, 1991, 1992). Such conditions are necessary for the forward problem, but they are not sufficient for the inverse problem. Even when all regularity conditions are satisfied and the transform is perfectly well defined, the model may still be structurally non-identifiable if the symmetry group is nontrivial. Regularity ensures that the model is a valid generative model; it does not guarantee uniqueness of parameter recovery.

When an OP diagram is structurally non-identifiable because $|\mathcal{G}(D)| > 1$, identifiability can be restored only by introducing additional information that breaks the symmetry. One possibility is an ordering constraint on rates, such as $\lambda_1 > \lambda_2$, which distinguishes otherwise exchangeable assignments. A second possibility is a selective-influence constraint, where an experimental manipulation is known to affect one process but not another, thereby separating structurally similar roles (Dzhafarov, 2003; Schweickert, 1985; Schweickert et al., 2000; Townsend & Liu, 2022). A third possibility is parameter linking, where theory entails equality or functional dependence among some rates, thereby reducing the effective parameter space. These devices do not alter the structural source of the ambiguity itself; rather, they supplement the observational model with theoretically justified constraints strong enough to remove it.

Two OP diagrams are said to belong to the same transform-equivalence class if they generate identical moment-generating functions after an appropriate relabelling of process-specific rates. Formally, D_1 and D_2 are equivalent if there exists a permutation σ such that

$$M_{T,1}(s \mid \lambda) = M_{T,2}(s \mid \sigma(\lambda)) \tag{2.26}$$

for all admissible s. This definition captures equivalence induced by relabelling within the present formalism. It is narrower than the broader notion of mimicry in the serial–parallel literature, where observational equivalence may obtain between models belonging to different architectural families (Townsend, 1972; Townsend & Ashby, 1983). Nevertheless, it is theoretically important, because it identifies when distinct-looking OP representations are not empirically distinguishable within the present transform-based framework.

In summary, the moment-generating-function approach yields two distinct levels of structural identifiability in the present subclass of OPM models. First, under the stated existence and generic no-cancellation conditions, the response-time transform identifies the multiset of rate parameters appearing as poles in the reduced rational form. Second, identifiability of the assignment of those rates to specific process labels depends on the symmetry structure of the diagram. Nontrivial symmetry groups imply assignment non-identifiability; trivial symmetry groups remove that specific source of ambiguity, though not every logically possible source of observational equivalence. This distinction between multiset recovery and assignment recovery provides the foundation for the later analysis of concrete OP structures.

3. Application to visual search models

Visual search provides a well-established domain for applying the identifiability framework. The task requires subjects to indicate as quickly as possible whether a target stimulus appears in a display of distractors. Goldstein and Fisher (1991) themselves illustrated the OPM using visual-search models, making this domain particularly relevant for evaluating the present analysis.

Consider first a simplified visual-search task with three identical display items. The assumed model is that all items are encoded in parallel and compared to the target representation in parallel. The network structure comprises three parallel paths, each representing an item, as illustrated in Fig. 2.

- Path 1: Encode item 1 (E_1) with rate λ_e, then compare with the target (C_1) with rate λ_c.
- Path 2: Encode item 2 (E_2) with rate λ_e, then compare with the target (C_2) with rate λ_c.

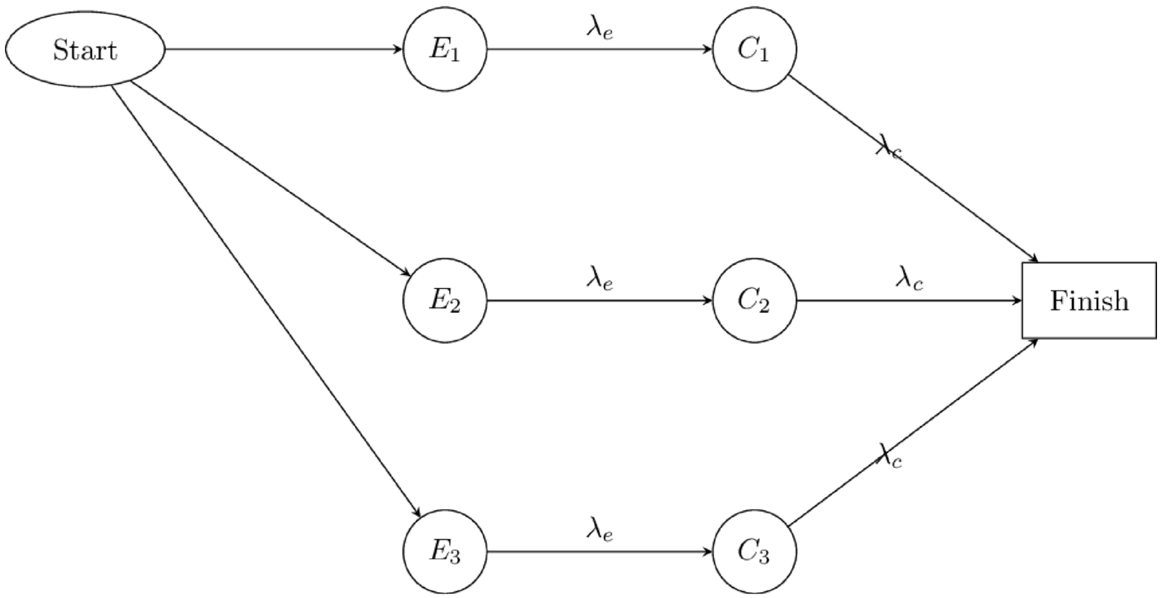

Fig. 2. Symmetric parallel Order-of-Processing (OP) diagram illustrating a model of visual search with three identical items. Each item is encoded (E_1, E_2, E_3) and compared (C_1, C_2, C_3) in parallel with shared rate parameters (λ_e for encoding and λ_c for comparison). The symmetry of this diagram induces assignment non-identifiability because multiple parameter assignments yield the same response-time distribution.

- Path 3: Encode item 3 (E_3) with rate λ_e, then compare with the target (C_3) with rate λ_c.

This architecture is symmetric because the three paths are structurally identical and exchangeable. Each item undergoes the same sequence of processes with the same path weight. The path-process incidence matrix for this model is

$$I = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}, \tag{3.1}$$

where columns represent processes E_1, C_1, E_2, C_2, E_3, and C_3, and rows represent paths. If the three paths are equiprobable, the response-time transform is

$$\begin{aligned} M_T(s) &= \frac{1}{3}\sum_{k=1}^{3} \frac{\lambda_e}{\lambda_e - s} \cdot \frac{\lambda_c}{\lambda_c - s} \\ &= \frac{\lambda_e \lambda_c}{(\lambda_e - s)(\lambda_c - s)}. \end{aligned} \tag{3.2}$$

By Theorem 2.4, any permutation that merely relabels the three exchangeable item paths leaves the transform unchanged. Consequently, item-specific assignment of encoding or comparison rates is not structurally identifiable in this model. What is recoverable is the common rate structure, not a unique mapping from those rates to particular item labels.

Consider next a modified visual-search model in which encoding is sequential but comparisons proceed in parallel as soon as their corresponding encodings finish. This is the hybrid serial–parallel model shown schematically in Fig. 3.

This architecture is asymmetric because the three encoding subprocesses occupy distinct serial positions, and the comparison subprocesses have staggered onset times. Let X_{e1}, X_{e2}, X_{e3} denote the durations of the encoding subprocesses and X_{c1}, X_{c2}, X_{c3} the durations of the comparison subprocesses. If the diagram in Fig. 3 is interpreted literally, then the overall response time is not a simple serial sum of all six durations. Rather, it is the time at which the latest of the three comparison branches finishes:

$$T_{\text{hyb}} = \max\left\{X_{e1} + X_{c1},\ X_{e1} + X_{e2} + X_{c2},\ X_{e1} + X_{e2} + X_{e3} + X_{c3}\right\}. \tag{3.3}$$

Accordingly, the transform of this hybrid model is *not* given by a simple product of exponential MGFs. That product form would be valid for a fully serial architecture, but the present diagram contains overlapping comparison branches and therefore involves a maximum over branch completion times.

If X_{e1}, X_{e2}, X_{e3} are independent exponentials with rates $\lambda_{e1}, \lambda_{e2}, \lambda_{e3}$, and X_{c1}, X_{c2}, X_{c3} are independent exponentials with common rate λ_c, then for fixed encoding durations u_1, u_2, u_3 the conditional distribution of T_{hyb} is

$$\begin{aligned} &P\left(T_{\text{hyb}} \le t \mid X_{e1} = u_1, X_{e2} = u_2, X_{e3} = u_3\right) \\ &\quad = \prod_{j=1}^{3}\left(1 - e^{-\lambda_c (t - s_j)}\right), \end{aligned} \tag{3.4}$$

whenever $t \ge s_3$, where

$$s_1 = u_1, \qquad s_2 = u_1 + u_2, \qquad s_3 = u_1 + u_2 + u_3.$$

Integrating over the encoding durations yields

$$\begin{aligned} F_{T_{\text{hyb}}}(t) = &\int_0^t \int_0^{t-u_1} \int_0^{t-u_1-u_2} \lambda_{e1} e^{-\lambda_{e1} u_1} \lambda_{e2} e^{-\lambda_{e2} u_2} \lambda_{e3} e^{-\lambda_{e3} u_3} \\ &\times \left(1 - e^{-\lambda_c (t-u_1)}\right)\left(1 - e^{-\lambda_c (t-u_1-u_2)}\right) \\ &\times \left(1 - e^{-\lambda_c (t-u_1-u_2-u_3)}\right) du_3\, du_2\, du_1. \end{aligned} \tag{3.5}$$

Eqs. (3.3)–(3.5) provide the correct completion-time law for the architecture as drawn. The important point for the present paper is conceptual as well as technical. This hybrid model lies outside the closed-form path-mixture subclass analysed in Section 2. Nevertheless, the diagram remains useful as a boundary example: because the encoding subprocesses occupy nonexchangeable serial positions, and because the comparison subprocesses have different effective onset times, the specific permutation symmetry responsible for the ambiguity in the symmetric example is absent here. This removes the particular symmetry-based non-identifiability studied in Section 2, although it does *not* by itself exclude broader mimicry by architectures outside the present subclass.

Finally, consider a partially symmetric visual-search model in which two item branches are exchangeable while a third occupies a distinctive structural role. Fig. 4 gives a schematic OP representation of this situation. The two red branches are structurally identical and therefore exchangeable; the blue branch is not. In such a case, the symmetry

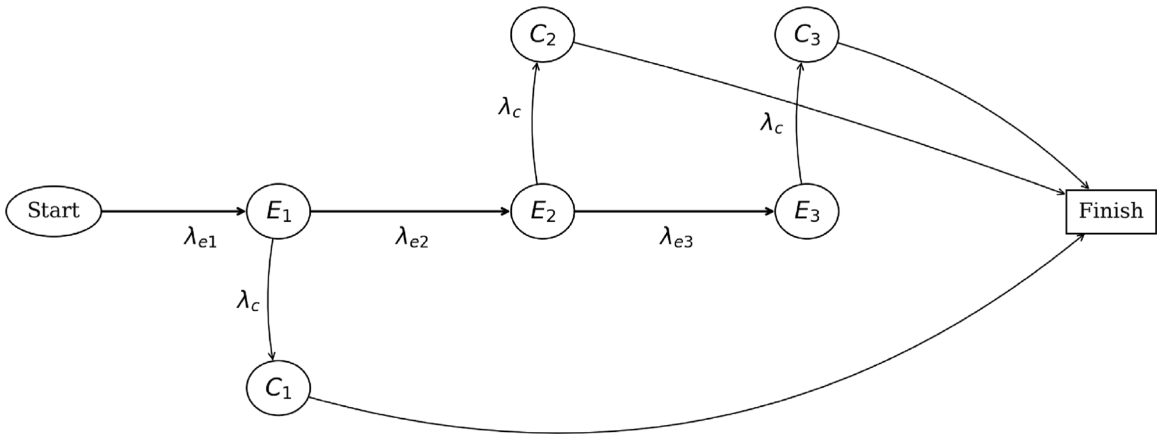

Fig. 3. Hybrid serial–parallel Order-of-Processing (OP) diagram for visual search. The encoding subprocesses E_1, E_2, and E_3 occur sequentially, while each comparison subprocess C_i begins as soon as the corresponding encoding subprocess finishes.

group contains the transposition exchanging the two red branches but excludes permutations involving the blue branch. Hence

$$\mathcal{G}(D) = \{id, \sigma_{12}\}, \qquad |\mathcal{G}(D)| = 2.$$

By Theorem 2.4, the model therefore retains a twofold ambiguity, but only with respect to the first two items. This residual ambiguity can be broken by an item-specific manipulation—for example, a stimulus-quality perturbation that selectively affects only one of the exchangeable branches.

4. Discussion

The theory developed here addresses a question in mathematical cognitive psychology: when can an OP model be interpreted as uniquely identifying psychologically meaningful rate parameters from response-time distributions, and when can it not? The main result is deliberately narrower than a universal theory of architectural identifiability. Within the restricted path-based subclass analysed here, the response-time transform generically identifies the multiset of rate parameters under the stated no-cancellation conditions, but assignment of those rates to specific process labels can fail whenever the OP diagram contains nontrivial permutation symmetries.

This point has practical implications. Before fitting an OPM to behavioural data and interpreting the resulting rate estimates psychologically, the researcher should first determine whether the proposed diagram has a nontrivial symmetry group. If it does, then the fitted parameters may be observationally valid without being uniquely attributable to specific process labels. The proper response is not to abandon the model automatically, but to add theoretically justified information strong enough to break the symmetry. Different symmetry types suggest different remedies. Exchangeable parallel branches call for branch-specific manipulations; exchangeable process families within a path call for selective-influence manipulations that affect one family but not another; and low-order residual ambiguities may be broken by theory-based ordering constraints or parameter-linking assumptions.

The relation to earlier selective-influence work is therefore substantive rather than merely historical. The present framework agrees with that literature that architectural inference depends not only on the observable data, but also on principled claims about what manipulations affect which subprocesses (Dzhafarov, 2003; Schweickert, 1985; Schweickert et al., 2000; Townsend & Liu, 2022). In this sense, selective influence is not merely a convenient interpretive device. It is one of the main ways in which permutation-based ambiguity can be reduced or removed.

At the same time, the present analysis should not be conflated with the classical serial–parallel mimicry problem. Townsend's results and the broader serial–parallel literature showed that distinct model classes can generate the same observable completion-time distributions (Townsend, 1972, 1990; Townsend & Ashby, 1983). The present paper makes a different point. Even *within* a fixed OPM subclass, ambiguity can remain because the diagram itself contains exchangeable structural roles. Thus, the paper complements rather than replaces the classical mimicry literature.

The logical independence of the two problems can be made precise through a concrete scenario. Consider a purely serial OPM with two processes A and B, each exponentially distributed with rates λ_a and λ_b respectively, and with a single fixed path on every trial. The response-time transform is

$$M_T(s) = \frac{\lambda_a}{\lambda_a - s} \cdot \frac{\lambda_b}{\lambda_b - s}, \tag{4.1}$$

and because multiplication is commutative, swapping the two rate labels leaves Eq. (4.1) unchanged. The symmetry group therefore contains the transposition $a \leftrightarrow b$, and by Theorem 2.4 the assignment of rates to process labels is not structurally identifiable: the multiset $\{\lambda_a, \lambda_b\}$ can be recovered, but not which rate belongs to which stage. Now consider the classical result that an exponential parallel model can, under appropriate conditions, generate the same marginal response-time distribution as this serial model (Townsend, 1972; Townsend & Ashby, 1983). That result concerns a different architecture, one in which both channels operate simultaneously and the overall response time is governed by a minimum or maximum completion time across channels. The parametric conditions under which a parallel model matches the serial distribution depend on the *shape* of the marginal distribution, not on which process label is assigned which rate within the serial model. Whether the serial rate assignment reads (λ_a, λ_b) or (λ_b, λ_a) is irrelevant to the question of whether a parallel architecture can also produce that same mixture of exponential terms. Conversely, an asymmetric serial OPM whose symmetry group is trivial — so that the within-class assignment non-identifiability is removed — can still be mimicked by a suitably parameterised parallel model, because mimicry depends on distributional equivalence across architectural families, not on uniqueness of parameter attribution within one family. The two inferential problems therefore operate at different levels: OPM permutation non-identifiability is a within-class ambiguity about label assignment given a fixed architecture; serial–parallel mimicry is a cross-class equivalence about architectural form given an observable distribution. Resolving one does not entail resolving the other.

The hybrid visual-search example in Section 3 is important in this regard. It shows that asymmetry in the diagram does not license careless simplification of the response-time law. Once comparison subprocesses overlap in time and finish at different branch-dependent

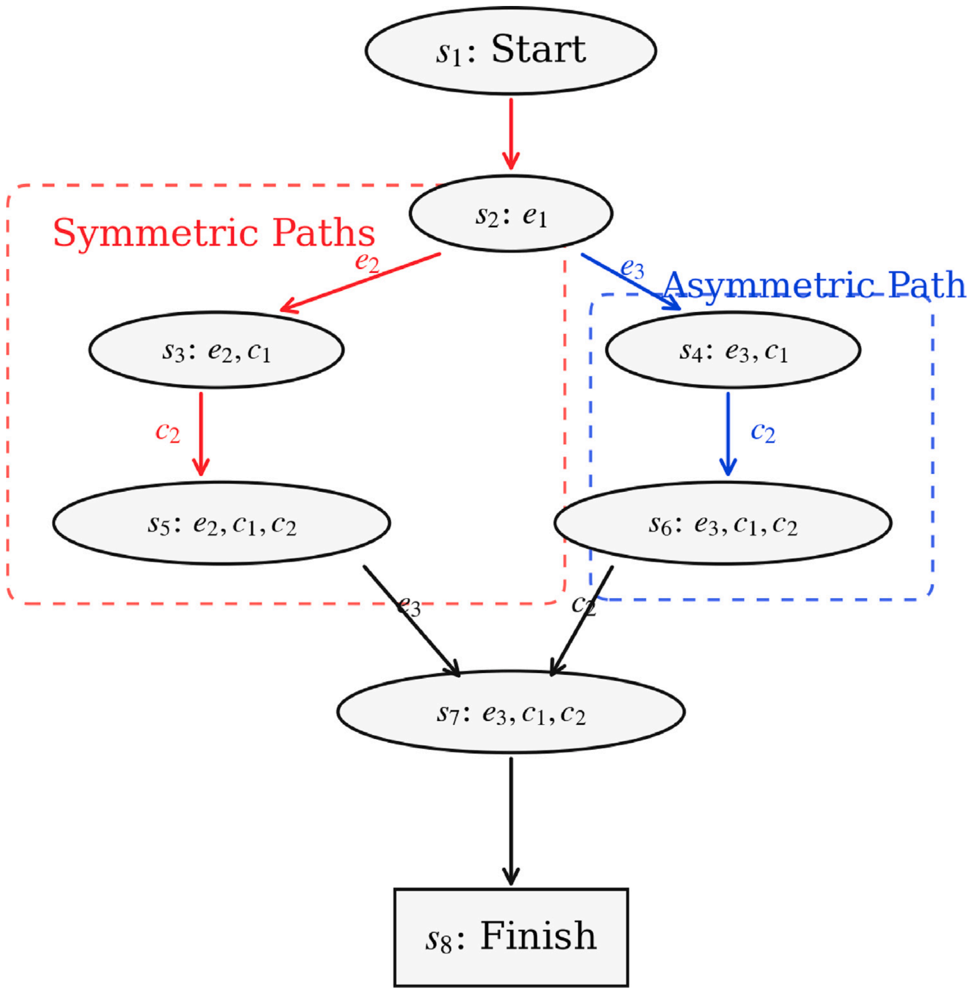

Fig. 4. Schematic Order-of-Processing (OP) diagram for a partially symmetric visual-search model. The red substructure contains two exchangeable branches, whereas the blue substructure contains a distinctive branch. The resulting symmetry group contains only the transposition of the two exchangeable branches.

completion times, the overall response time is governed by a maximum, not by a simple serial sum. That example therefore serves a dual role. It illustrates how symmetry may be absent while transform derivation remains mathematically nontrivial, and it clarifies the boundary of the closed-form path-mixture subclass studied in Section 2. This is why the present theory is framed as a restricted structural-identifiability analysis rather than as a universal theorem about all OP architectures.

The computation of the symmetry group is itself algorithmic. Starting from the path-process incidence structure and path weights, one can construct a coloured bipartite graph and compute its automorphism group using standard graph-automorphism software. This provides a practical route for applying the theory in concrete model-building contexts. Although we are not aware of a dedicated software package designed specifically for OPM identifiability, the underlying computational problem is standard enough that existing graph-automorphism tools are sufficient.

The present analysis is also restricted in distributional scope. Exponential process durations were adopted because they yield explicit transforms and transparent pole structure. Extensions to other distribution families, including gamma and Weibull forms, would be natural, but would require new transform characterisations and careful attention to whether analogous pole-based arguments remain available. Likewise, the present analysis presumes a finite OP diagram with a countable path set. More complex diagrams, continuous state spaces, or dynamically evolving path structures would require additional development.

A second limitation concerns the broader space of cognitive dynamical systems. The present results apply to discrete random-duration models of the kind represented by OP diagrams. They do not directly extend to continuous-flow or cascade conceptions in which downstream processing can begin before upstream processing has completed, or to architectures more naturally formulated through systems of differential equations (McClelland, 1979). Such models raise different identifiability questions because the relevant observables are no longer determined by pathwise completion alone. The present paper therefore clarifies one important inferential problem, but not every inferential problem that arises in mathematical cognitive modelling.

Finally, the paper concerns structural identifiability rather than practical identifiability. Structural identifiability asks whether different parameter assignments can in principle generate the same observable distribution. Practical identifiability asks whether, with finite and noisy data, parameters can be estimated with acceptable precision. A diagram may be structurally identifiable yet practically unstable if the observable distinctions among nearby parameter values are too small. Conversely, a symmetry-induced non-identifiability cannot be repaired by more data alone. The distinction between these two levels of inference remains important for future work.

5. Conclusion

The Order-of-Processing framework introduced by Goldstein and Fisher provides a mathematically elegant method for deriving response-time distributions from cognitive network models. The forward problem—deriving observable distributions from assumptions about structure and process durations—is well understood. The present paper addresses a narrower inverse problem: within a restricted path-based subclass of OPM models, when does the response-time distribution determine the underlying rate parameters, and when does it not?

The principal result is that structural ambiguity in this subclass arises naturally from permutation symmetry in the OP diagram. Under the stated existence and no-cancellation conditions, the reduced rational transform generically identifies the multiset of rate parameters. However, assignment of those rates to specific process labels can fail whenever nontrivial permutations preserve the response-time transform. The symmetry group of the diagram provides an explicit mathematical representation of this ambiguity and a practical guide to its severity.

This conclusion has direct methodological consequences. Before fitted OPM parameters are interpreted psychologically, the proposed diagram should be checked for symmetry-based non-identifiability. If such ambiguity is present, it must be broken by additional theory, experimentally induced asymmetry, selective-influence manipulations, or parameter constraints. In this way, the paper contributes not a universal solution to all architecture-identification problems, but a rigorous clarification of one important inferential limit within the OPM tradition.

Important limitations remain. The results apply to a restricted path-based subclass, not to the whole OPM and not to all cognitive architectures. The visual-search applications show both the usefulness of the present framework and the point at which more general overlapping-process architectures require a different mathematical treatment. Nevertheless, by clarifying when response-time distributions do and do not support unique parameter attribution, the present paper strengthens the inferential discipline with which OPM-based models can be used.

Funding

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 101034319 and from the European Union – NextGenerationEU.

Acknowledgments

The author would like to express sincere gratitude to reviewers for their constructive feedback and comments, which have improved the quality of this work and contributed to its present form.

Data availability

No data was used for the research described in the article.

References

Curtiss, J. H. (1942). A note on the theory of moment generating functions. *The Annals of Mathematical Statistics*, *13*(4), 430–433. http://dx.doi.org/10.1214/aoms/1177731541.

Donders, F. C. (1868). Die schnelligkeit psychischer processe: Erster artikel. *Archiv Für Anatomie, Physiologie Und Wissenschaftliche Medicin*, 657–681.

Dzhafarov, E. N. (2003). Thurstonian-type representations for "same-different" discriminations: Deterministic decisions and independent images. http://dx.doi.org/10.1016/S0022-2496(02)00027-5.

Goldstein, W. M., & Fisher, D. L. (1991). Stochastic networks as models of cognition: Derivation of response time distributions using the order-of-processing method. *Journal of Mathematical Psychology*, *35*(2), 214–241. http://dx.doi.org/10.1016/0022-2496(91)90026-P.

Goldstein, W. M., & Fisher, D. L. (1992). Stochastic networks as models of cognition: Deriving predictions for resource-constrained mental processing. http://dx.doi.org/10.1016/0022-2496(92)90055-C.

McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*(4), 287–330. http://dx.doi.org/10.1037/0033-295X.86.4.287.

Schweickert, R. (1978). A critical path generalization of the additive factor method: Analysis of a stroop task. *Journal of Mathematical Psychology*, *18*(2), 105–139.

Schweickert, R. (1980). Critical-path scheduling of mental processes in a dual task. *Science*, *209*(4457), 704–706. http://dx.doi.org/10.1126/science.7394529.

Schweickert, R. (1982). The bias of an estimate of coupled slack in stochastic PERT networks. *Journal of Mathematical Psychology*, *26*(1), 1–12.

Schweickert, R. (1985). Separable effects of factors on speed and accuracy: Memory scanning, lexical decision, and choice tasks. http://dx.doi.org/10.1037/0033-2909.97.3.530.

Schweickert, R. (1992). Constructing a stochastic critical path network given the slacks: Representation. *Mathematical Social Sciences*, *23*(3), 343–366. http://dx.doi.org/10.1016/0165-4896(92)90046-8.

Schweickert, R., Giorgini, M., & Dzhafarov, E. (2000). Selective influence and response time cumulative distribution functions in serial-parallel task networks. *Journal of Mathematical Psychology*, *44*(4), 504–535. http://dx.doi.org/10.1006/jmps.1999.1268.

Schweickert, R., & Townsend, J. T. (1989). A trichotomy: Interactions of factors prolonging sequential and concurrent mental processes in stochastic discrete mental (PERT) networks. *Journal of Mathematical Psychology*, *33*(3), 328–347. http://dx.doi.org/10.1016/0022-2496(89)90013-8.

Singhal, K., & Vlach, J. (1975). Computation of time domain response by numerical inversion of the Laplace transform. *Journal of the Franklin Institute*, *299*(2), 109–126. http://dx.doi.org/10.1016/0016-0032(75)90133-7.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. *Acta Psychologica*, *30*, 276–315. http://dx.doi.org/10.1016/0001-6918(69)90055-9.

Sternberg, S. (1984). Stage models of mental processing and the additive-factor method. *Behavioral and Brain Sciences*, *7*(1), 82–84. http://dx.doi.org/10.1017/S0140525X00026285.

Stoyanov, J. (2015). Moment properties of probability distributions used in stochastic financial models. In *Recent advances in financial engineering 2014* (pp. 1–27). World Scientific, http://dx.doi.org/10.1142/9789814730778_0001.

Townsend, J. T. (1972). Some results concerning the identifiability of parallel and serial processes. *British Journal of Mathematical and Statistical Psychology*, *25*(2), 168–199. http://dx.doi.org/10.1111/j.2044-8317.1972.tb00490.x.

Townsend, J. T. (1984). Uncovering mental processes with factorial experiments. *Journal of Mathematical Psychology*, *28*(4), 363–400. http://dx.doi.org/10.1016/0022-2496(84)90007-5.

Townsend, J. T. (1990). *It was twenty years ago today: Increasing reaction times or errors still do not a serial processor make: Tech. rep.*, West Lafayette, IN: Purdue University, Purdue Mathematical Psychology Program.

Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. CUP Archive: Accessed 08 May 2025.

Townsend, J. T., & Liu, Y. (2022). Varieties of selective influence: Toward a more complete taxonomy and implications for systems identification. *Mathematics*, *10*(7), 1059. http://dx.doi.org/10.3390/math10071059.

Townsend, J. T., & Schweickert, R. (1989). Toward the trichotomy method of reaction times: Laying the foundation of stochastic mental networks. *Journal of Mathematical Psychology*, *33*(3), 309–327.