

State-Task Alignment for Consequence-Sensitive Task Routing

A Three-Layer Research Framework Under Uncertain Readiness

James McGaughran · ORCID 0009-0005-3324-7290

MathGov / RippleLogic

Working Paper v1.6 · August 2026

Draft status. This manuscript proposes a falsifiable integrative framework. It is not peer reviewed, is not a validated assessment or scheduling instrument, and supplies no clinical, legal, organizational, or safety certification. A reproducible scoping review and external specialist review are required before journal submission.

Version note. Version 1.6 is the convergence draft. It demotes sharp change-point recovery from the confirmatory core, makes measurement quality the empirical go/no-go condition, removes unreproducible simulation claims from the manuscript, repairs the reference audit, adds operator-performance prediction evidence, and sharpens the EU worker-allocation boundary. The framework now advances four confirmatory claims and one explicitly exploratory functional-form question. A final independent verification pass corrected the last stale hypothesis label, added the currently applicable EU transparency and AI-literacy duties, and hardened the citation audit against same-surname, same-year collisions.

Abstract

Work-allocation systems commonly prioritize tasks by importance, urgency, role, or clock time while representing the worker's present functional readiness only weakly. Research on temporal person-environment fit, operator functional state, sleep, fatigue, affect, multiple resources, recovery, adaptive allocation, and contingent planning shows that both human state and effective task demand vary. Safety engineering further shows that safeguards can determine whether an activity may proceed. These traditions do not by themselves provide a rights-constrained policy for selecting among a person's available task episodes when readiness is uncertain and consequences, detectability, reversibility, and safeguards differ.

This paper develops **State-Task Alignment (STA)** as a three-layer research framework. The **descriptive layer** tests task-relative relations among readiness evidence, demand, controls, and outcomes. The **prescriptive layer** compares task-control bundles and permits select, select with controls, prepare only, defer, restore, care/non-work, escalate, or declare structural constraint. The **normative layer**, derived from MathGov/RippleLogic's qualify-before-rank architecture, places rights, severe risk, and structural viability outside compensatory optimization. These layers are intentionally distinct: empirical evidence can test readiness-outcome and policy effects, but it cannot by itself validate a rights floor.

Successive simulation audits showed that recovering a sharp change-point from noisy readiness evidence is highly sensitive to measurement error, null-family choice, calibration

stability, and person-varying cutoffs. Those diagnostics are exploratory design work, not evidence that thresholds exist. Version 1.6 therefore removes sharp threshold recovery from the confirmatory core. STA's routing logic requires only that task-relative readiness and demand predict outcomes sufficiently to support a decision rule; any operational cutoff must be justified by validated evidence, professional standards, consequence, or precaution rather than presented as a discovered biological boundary.

The confirmatory program now asks four questions: whether task-relative readiness-demand correspondence improves objective prediction within persons; whether specified safeguards change identifiable mechanisms; whether the full routing policy outperforms a strong simple contingent heuristic; and whether it remains within primary safety, autonomy, equity, and burden margins. Functional form, including possible change-points and insight-task reversals, remains exploratory until Phase 2 produces sufficiently reliable task-family-specific measures. Until a reproducible scoping review, measurement program, and comparative trials are completed, STA is a structured hypothesis and auditable research design, not validated guidance.

Keywords: within-person fit; operator functional state; functional readiness; task allocation; adaptive intervention; human factors; safeguards; irreversibility; self-regulation; rights-constrained optimization

Claim boundary. Established evidence supports within-person variation in performance-relevant states, heterogeneous and task-conditional timing effects, dissociable task demands, domain-specific impairment from sleep loss and fatigue, imperfect self-assessment, switching costs, contingent-plan effects, and limited recovery effects. It does **not** show that the integrated STA policy improves real-world outcomes, that a universal readiness ontology exists, that a natural threshold can be recovered from brief self-ratings, or that organizational or automated deployment is safe. Simulation can expose design weakness; it cannot validate the theory.

1. Introduction: From Time Allocation to Task-Episode Routing

A scheduled hour is not a fungible unit of human performance. The same person can be analytically clear but emotionally reactive, physically able but precision-poor, creatively open but unable to sustain controlled checking, or willing to work while objectively impaired. Tasks are equally heterogeneous. Drafting a theory, reconciling immigration declarations, clearing a workspace, conducting a difficult conversation, and recovering from accumulated strain demand different capabilities and carry different consequences of error or omission.

This creates a recurrent routing problem: given a bounded work interval, a present state known only imperfectly, and several candidate tasks, which task — with which safeguards — should receive the interval? Priority systems answer mainly from the task side. Chronotype and fatigue advice answer mainly from the person side. STA treats allocation as a joint, dynamic, and consequence-sensitive problem.

The initiating intuition is simple: direct strong states toward work that needs them, and use depleted states for safe maintenance or restoration. That intuition is useful but scientifically incomplete, and each of its failures is instructive.

- Tedious administrative work may be high-demand and high-consequence.
- Some insight tasks improve when inhibitory control is *reduced*, so “more readiness is better” is false for part of the task space.
- Low-capacity maintenance can either prepare the next work episode or consume needed recovery.
- Felt readiness diverges from performance, and diverges in the dangerous direction.
- Urgent work may require stronger controls rather than faster finalization.

The paper therefore replaces a high-energy/low-energy slogan with an explicit qualification and routing architecture.

The unit of analysis is the **task episode**: a bounded task component performed under a specified control bundle in a particular context. The theory does not assume that a whole project has one demand profile. A legal application decomposes into low-risk file preparation, moderate-risk transfer of verified data, high-risk interpretation, and an irreversible final submission — four episodes with four different eligibility conditions.

One structural commitment should be visible from the outset. Compensatory ranking can allow a sufficiently valuable option to outrank a poorly matched one. That property is unacceptable for consequential work when a designated condition is not met. STA therefore qualifies options before ranking them. The non-compensatory decision cutoff may be empirical, professional, legal, or precautionary; the framework no longer assumes that it corresponds to a sharp natural threshold in human performance.

2. Positioning, Lineage, and Bounded Contribution

2.1 Temporal fit and demands-abilities fit

Classic person-environment and demands-abilities fit research established that outcomes partly depend on correspondence between capabilities and demands measured on commensurate dimensions (Kristof-Brown, Zimmerman, & Johnson, 2005). Temporal fit research explicitly studies how fit changes within people, situations, and time (Vleugels, Verbruggen, De Cooman, & Billsberry, 2023). STA therefore cannot claim that momentary fit is an unoccupied problem.

Fit research has also tested nonlinearity. Vleugels and Flatau-Harrison (2023) reported curvilinear relations for some behavioral outcomes in a two-wave employee study. This defeats any generic novelty claim about curvilinear fit, but it is not a direct test of momentary task episodes, binding readiness, control bundles, or critical errors. Response-surface and polynomial-regression methods further caution against reducing commensurate person and task profiles to one difference score (Edwards & Parry, 1993). STA must test readiness, demand, context, and controls directly.

2.2 Operator functional state and adaptive function allocation

Operator-functional-state research is a closer incumbent than earlier drafts acknowledged. It studies changing workload, fatigue, stress, vigilance, and related states in supervisory and safety-sensitive work, frequently using physiological and behavioral signals. A recent scoping review mapped 109 empirical studies and found a heterogeneous evidence base, with considerably more support for detecting mental-state correlates than for reliably predicting operational performance (Pütz, Mertens, Chuang, & Nitsch, 2024). A follow-up study found that pupil size and heart-rate variability can contribute to operator-performance prediction, but that confounding changes in the effort needed to maintain performance limit simple state-to-performance inference (Pütz, Mertens, Chuang, & Nitsch, 2025). Earlier work explicitly proposed dynamically allocating operator task load according to fluctuating psychophysiological functional state (Zhang et al., 2017).

This lineage sharply limits STA's originality. Measuring state, adapting task load, and using multidimensional operator models are not new. STA's extension is the voluntary routing of a portfolio of ordinary cognitive, administrative, relational, and maintenance episodes; separation of omission loss from execution risk; explicit uncertainty about self-estimation; safeguard-dependent admissibility; protected non-work outputs; and a rights floor above the policy objective. Whether that transfer adds empirical value is a research question, not an established contribution.

2.3 Multiple resources, fatigue, and effort allocation

Multiple resource theory models differentiated resource demands and interference, especially in concurrent performance (Wickens, 2002, 2008). Motivational accounts interpret fatigue as a control signal involving changing costs, goals, and effort allocation rather than a literal empty fuel tank (Boksem & Tops, 2008; Hockey, 2013). Expected-value-of-control, opportunity-cost, labor/leisure, and resource-rational approaches model whether limited control should be invested in a task (Kurzban et al., 2013; Shenhav, Botvinick, & Cohen, 2013; Kool & Botvinick, 2014; Lieder & Griffiths, 2020).

STA borrows these foundations and asks a portfolio question: among the available task-control bundles, which admissible episode should receive the next bounded interval?

2.4 Adaptive interventions, automation, and safety controls

Human-factors research has dynamically reallocated functions between people and automation using operator state, task performance, and workload (Parasuraman, Mouloua, & Molloy, 1996; Kaber & Endsley, 2004). Just-in-time adaptive intervention research formalizes decision points, tailoring variables, decision rules, and proximal outcomes (Klasnja et al., 2015; Nahum-Shani et al., 2018). Permit-to-work systems and barrier engineering condition authorization on specified controls and residual risk (Health and Safety Executive, n.d.; Mancuso et al., 2017). STA therefore does not invent adaptive allocation or the control bundle as an authorization object.

Emerging human-AI allocation work also overlaps directly. The 2026 HAAS preprint combines governance constraints applied before learning with contextual-bandit allocation across five cognitive task dimensions (Pelechano, Mestre, Albert, & Gil, 2026). Because it is a recent preprint and concerns human-AI collaboration rather than voluntary self-routing, it is treated

as an emerging parallel rather than settled precedent. It nonetheless confirms that governance-before-optimization and auditable multidimensional allocation are active research territory. Detailed comparison belongs in the companion AI-governance paper.

2.5 Prior use of the expression

Dust (2015) used the expression *state-task alignment* in a commentary about matching mindfulness, flow, and mind-wandering states to task requirements. The present paper does not claim to coin the term. Dust's discussion concerned states of consciousness rather than the present demand-risk-control-governance architecture, but priority must be acknowledged precisely.

2.6 Defensible contribution

Proposed contribution. STA is an integrative extension that separates descriptive readiness-demand science, prescriptive task-control routing, and normative qualification; treats self-estimates as uncertain evidence; permits safeguards to alter admissibility; represents non-selection and structural redesign as legitimate outputs; and tests the complete policy against simpler contingent rules under primary safety constraints.

The contribution is architectural and translational, not the invention of its components. Four confirmatory questions carry the scientific burden:

1. Does task-relative readiness-demand correspondence predict objective episode outcomes within persons?
2. Do specified controls change identifiable mechanisms rather than merely reassure the user?
3. Does the full policy outperform a simple contingent heuristic at acceptable burden?
4. Does the policy remain non-inferior on sleep, avoidance, autonomy, equity, distress, and relational harm?

A fifth question - whether some task families show recoverable change-points or other distinctive functional forms - is exploratory and conditional on adequate measurement.

3. Scientific Foundations and Claim Limits

3.1 Time of day and chronotype

A 2025 systematic review located 65 studies. More than 80% reported no main effect of chronotype; among studies of adults aged 18-45, synchrony appeared in 29 of 64 comparisons, concentrated in attention, inhibition, and memory. Older-adult studies more often reported synchrony, including fluid-ability tasks (Chauhan et al., 2025), consistent with earlier evidence of synchrony effects in inhibitory control among older adults (May, Hasher, & Stoltzfus, 1998). A young-adult inhibition study found no synchrony effect (Tseng & Damian, 2023). The implication is not that timing is irrelevant, but that clock time and chronotype are weak substitutes for direct, task-specific state evidence.

Some task relationships reverse. Wieth and Zacks (2011) found insight problems performed better at non-optimal times, plausibly because reduced inhibitory filtering broadened associative access. STA must therefore reject the slogan that more control or more activation is always better.

3.2 Sleep loss and self-assessment

Sleep restriction and deprivation impair vigilance, executive functions, working-memory maintenance, and cognitive flexibility, with substantial heterogeneity across domains and individuals (Van Dongen et al., 2003; Lowe et al., 2017; Cao, Xie, & Ma, 2025). Crucially, subjective sleepiness can adapt while objective deficits continue accumulating, and fatigue-risk modeling shows that self-reported sleepiness can substantially understate accumulating performance impairment in operational schedules (McCauley et al., 2021). Self-report is evidence, not authorization. High-consequence routing requires stronger external evidence or controls when sleep debt or calibration failure is plausible.

3.3 Mental fatigue, effort, and vigilance

Mental fatigue can alter effort valuation, persistence, and performance, but subjective fatigue and objective impairment do not map one-to-one. Daylong cognitive work has been linked to altered economic choice, and causal work implicates lateral prefrontal processes in effort and fatigue judgments, although neither establishes a general routing rule (Blain, Hollard, & Pessiglione, 2016; Soutschek & Tobler, 2020; Boksem & Tops, 2008; Hockey, 2013; Jurgelis et al., 2021). Vigilance work is experienced as stressful and resource demanding even when overt motor demands are low (Warm, Parasuraman, & Matthews, 2008). STA therefore separates readiness, willingness, strain, and task demand instead of treating them as one energy scalar.

3.4 Affect, activation, and creativity

Affect-creativity relations are conditional. Positive activating states show modest average advantages, while negative affect can support creativity under some motivational and contextual conditions (Baas, De Dreu, & Nijstad, 2008; George & Zhou, 2002). Activation and associative breadth are therefore modifiers with task-dependent optima, not universally beneficial capacities.

3.5 Recovery, maintenance, and non-work

Micro-breaks show small average improvements in vigor and fatigue, but aggregate performance effects are uncertain and benefits are weaker for demanding cognitive tasks; duration and activity also matter (Albulescu et al., 2022; Bennett, Gabriel, & Calderwood, 2020). Recovery experiences include detachment, relaxation, mastery, and control (Sonnentag & Fritz, 2007), and effort-recovery research treats recovery as regulation across time rather than a simple refill (Meijman & Mulder, 1998; Zijlstra et al., 2014). These literatures do not authorize the claim that every low-readiness period should be filled with maintenance. Rest, care, leisure, and non-work may be intrinsically protected rather than justified by later output.

3.6 Contingent planning, switching, and procrastination

Implementation intentions improve goal pursuit on average, with recent synthesis showing stronger effects for contingent if-then plans than for scheduling-only plans (Gollwitzer & Sheeran, 2006; Sheeran, Listrom, & Gollwitzer, 2025). Experimental work also indicates that implementation intentions can modulate switching and interference performance (Cohen, Bayer, Jaudas, & Gollwitzer, 2008). These findings support the plausibility of contingent routing rules, not STA itself. Task switching carries measurable time and accuracy costs, so over-responsive rerouting can erase matching gains (Rubinstein, Meyer, & Evans, 2001; Monsell, 2003).

Procrastination research identifies task aversiveness, delay, low self-efficacy, and impulsiveness as important predictors (Steel, 2007), while temporal motivation theory models value, expectancy, delay, and impulsiveness jointly (Steel & König, 2006). These provide the necessary counterweight to readiness-based deferral.

3.7 Cautionary failures

Ego depletion is a cautionary precedent: an intuitive capacity theory achieved broad adoption before preregistered multi-laboratory tests substantially weakened its canonical effect (Inzlicht & Schmeichel, 2012; Hagger et al., 2016; Vohs et al., 2021). Decision-fatigue claims have also proved sensitive to design and confounding; reanalysis of the well-known judicial-sequence example showed that the apparent magnitude can be produced or inflated by case-ordering processes (Glöckner, 2016). STA should therefore privilege preregistered active-control tests, design analysis, and adverse-outcome measurement over plausibility.

3.8 What is not established

No direct evidence shows that momentary readiness-demand routing improves objective outcomes. No validated STA instrument exists. No universal readiness dimensions or thresholds have been established. No evidence supports clinical, educational, organizational, or automated deployment. Existing research supports premises and mechanisms; the integrated policy remains untested.

4. Three-Layer Construct Architecture

STA previously mixed descriptive constructs, action rules, and moral constraints in one architecture. Version 1.6 separates them because each has a different kind of warrant and a different failure mode.

4.1 Layer A - descriptive science

The descriptive layer asks what predicts episode outcomes. Its variables are empirical and revisable.

Table 2. Person-side objects.

Object	Provisional content	Scientific status
Task-relative readiness evidence	Integrative clarity; executive control; regulation stability; precision/error monitoring; physical readiness	Candidate domains, not a validated universal ontology. Their number and separability must be tested within persons and by task family.
State modifiers	Activation/arousal; affective valence; associative breadth	Task-conditional modifiers with possible inverted-U or reversed relations. They are not globally ordered capacities.
Willingness and valuation	Motivation; effort willingness; aversion; perceived reward	Choice variables, separated from capability so dislike cannot automatically become incapacity.
Evidence profile	Sleep and context; recent task-relevant performance; calibration history; source quality; uncertainty	Determines claim authority and required safeguards; never converted into a universal confidence score without calibration.

The five readiness domains are placeholders for measurement development. A general-factor solution, task-family-specific factors, or fewer dimensions would narrow the theory rather than falsify the routing problem. No domain should be inferred covertly from physiology for organizational use at the present evidence level.

Table 3. Task-side empirical layers.

Layer	Variables	Question
Demand	Person-relative cognitive, executive, regulatory, precision, physical, social, or associative requirements	What must this person deploy for this episode?
Execution risk	Error probability, impact, detectability, reversibility, propagation, affected parties	What happens if it is done badly now?
Value	Expected benefit, avoided omission loss, developmental and option value	What is gained by completion or lost by omission?
Constraints	Deadline, minimum block, setup/switching, sequence, dependencies, coordination, discretion	What limits feasibility?
Controls	Checklist, source trace,	Which causal mechanism

Layer	Variables	Question
	second reader, delayed finalization, staged action, scope reduction, delegation	can be changed?

Demand is person-relative, risk is not value, and controls must be identified by mechanism. A second reader improves detection, not the first actor's readiness. A template may reduce effective demand. Delayed finalization increases reversibility. Delegation changes the actor and authority boundary.

4.2 Layer B - prescriptive routing policy

The prescriptive layer maps declared evidence and an option set to actions. It is tested by comparative policy performance, not by whether every descriptive construct is novel. Valid outputs are: **select**, **select with controls**, **prepare only**, **defer to a reserved window**, **restore**, **care/non-work**, **escalate/redesign**, and **declare structural constraint**.

A policy can be useful even if readiness-outcome relations are smooth rather than thresholded. Conversely, a descriptive threshold can be real while a complex routing protocol adds no value beyond a simple rule. These claims must be tested separately.

4.3 Layer C - normative qualification

The normative layer states what the policy is not permitted to trade away: health, sleep, dignity, privacy, non-coercion, accommodation, contestability, protected care and non-work, and authority limits for consequential actions. These constraints are reasons-governed commitments supplied by the MathGov/RippleLogic source architecture and relevant law and professional norms. Empirical studies can test their feasibility, burden, distributional effects, and failure modes; they cannot prove the rights floor true by showing productivity gains.

4.4 Functional annotations

Generative, protective, maintenance, restorative, relational, and developmental remain non-exclusive annotations. They support decomposition and communication but carry no independent routing authority. The empirical layers, policy action set, and normative constraints do the work.

5. Formal Decision Architecture

The formalism is a preregistration scaffold, not a calculator. It deliberately separates empirical prediction, policy choice, and normative constraints.

5.1 Descriptive outcome model

For person p , episode i , control bundle m , and decision point t , let:

- R_{pt} : task-relevant readiness evidence and modifiers;
- $D_{pi}(m)$: person-relative demand after controls;

- X_pit: skill, sleep, time, task identity, environment, and other context;
- Y_pitm: objective episode outcome.

The descriptive claim is estimated as an outcome surface:

$$Y_{pitm} = f_{task}(R_{pt}, D_{pi}(m), X_{pit}) + \text{person/task effects} + \text{error}$$

No universal scalar fit score is required. Linear, response-surface, spline, change-point, and task-specific models are competitors. Functional form is exploratory until task-family-specific measurement is sufficiently reliable. A **binding condition** in the policy is a declared non-compensatory decision requirement; it need not be an empirically sharp change-point. Where a natural change-point is proposed, it must outperform smooth alternatives on independent data and its location must be sufficiently recoverable.

5.2 Control mechanisms

Each control bundle is represented by its proposed mechanism:

- demand reduction or redistribution;
- error prevention;
- increased detectability;
- increased reversibility or containment;
- reduced scope or authority;
- delegation to another qualified actor;
- burden in time, coordination, privacy, and opportunity.

A control claim is confirmed only for the mechanism and task family tested. “Use a checklist” is not a general theory of admissibility.

5.3 Qualification and action set

For each task-control bundle, the routing policy first applies non-compensatory qualification:

QUALIFY(i,m,t):

evidence authority is adequate for the proposed claim;
rights and authority conditions pass;
structural viability passes or passes with controls;
residual severe risk is below the declared tolerance;
any empirically or professionally justified binding condition is met.

Where no calibrated binding threshold exists, the policy may still require a control because of professional standards, consequence, low detectability, or low reversibility. It must label that rule as normative or precautionary rather than presenting it as an empirically discovered readiness threshold.

5.4 Ranking survivors

Only qualified bundles enter compensatory comparison:

PolicyValue = expected benefit + avoided omission loss
 - residual execution loss - control burden
 - setup/switching cost - expected downstream costs

Early studies should report outcome vectors and Pareto trade-offs rather than impose universal weights. The rights floor and severe-risk conditions remain outside the value function.

5.5 Valid outputs

Output	Meaning
SELECT	Admissible and preferred without added controls.
SELECT WITH CONTROLS	Admissible only with a specified safeguard bundle.
PREPARE ONLY	Reversible preparation is admissible; finalization is not.
DEFER TO RESERVED WINDOW	A later window is plausible, delay is acceptable, and the reservation is recorded.
RESTORE	Recovery is warranted or continued productive work is inadmissible.
CARE / NON-WORK	Sleep, care, relationship obligations, leisure, or stopping are protected without productivity justification.
ESCALATE / REDESIGN	Obtain help, change scope, renegotiate, redesign, or delegate.
DECLARE STRUCTURAL CONSTRAINT	Staffing, scheduling, authority, tools, or the option set are inadequate.

5.6 What survives if thresholds fail

Sharp threshold structure is not the identity of STA. The descriptive layer may be graded, curved, task-specific, or discontinuous. A policy cutoff may be set from validated predictive risk, professional standards, legal authority, consequence, detectability, reversibility, or precaution. Control effectiveness, protected non-selection, and comparative routing remain separately testable. If the full policy does not outperform the simple heuristic, the public protocol must be simplified.

6. Operating Protocol

Table 6. READ – GATE – MATCH – RANK – COMMIT – REVIEW.

Stage	Decision	Required output
READ	Estimate readiness, modifiers, willingness,	A state estimate with declared uncertainty, not a

Stage	Decision	Required output
	context, and evidence quality for the next block	global energy label
GATE	Check rights, safety duties, binding demands, tail risk, reversibility, detectability, structural feasibility	Eligible task–control bundles, or a justified empty set
MATCH	Identify binding dimensions and the controls that change demand or risk	A small set of defensible episode–control pairs
RANK	Compare benefit, avoided omission loss, expected error loss, burden, switching, future-capacity effects	Selected qualified bundle, and the reason
COMMIT	Work for a bounded interval without continuous rerouting	Protected block, scope, stopping condition, finalization limit
REVIEW	Compare forecast and outcome; update calibration and backlog	Outcome record, error and rework data, state change, next action

A practical public heuristic can be stated without presenting the theory as validated:

When clear, do what requires clarity. When steady, do what requires steadiness.
When depleted, use safe maintenance, restore, care, or stop. When consequences are high, do not confuse urgency with readiness.

The empirical question is whether the full architecture outperforms this rule by enough to justify its burden. That is not a rhetorical flourish; it is confirmatory claim C3, and the answer may be no.

7. Anti-Avoidance and Protected Restoration

A state-responsive framework becomes a sophisticated **avoidance** system unless deferral is observable and costly. It becomes a sophisticated **extraction** system unless health, leisure, care, and non-work time are protected. STA needs both constraints, and they pull in opposite directions, which is why each is stated explicitly rather than left to judgement.

Table 7. Anti-avoidance and anti-extraction rules.

Rule	Function
Sufficient, not ideal	A task needs an adequate state and control bundle, not a perfect mood
Next adequate window	Important deferred tasks are reserved for the

Rule	Function
	next episode meeting the declared binding conditions, and the reservation is recorded
Count deferral	Repeated postponement is data; it triggers the tree in Appendix D
Separate unwillingness from unreadiness	Aversion, fear, ambiguity, skill deficit, moral refusal, and genuine incapacity require different responses
Bounded initiation	A short reversible start can test a pessimistic estimate — but never overrides a rights or safety gate
Urgency changes controls, not reality	A deadline may justify assistance, scope reduction, or staged submission; it does not manufacture readiness
Rest is not only instrumental	Restoration may be intrinsically justified and may have delayed or unmeasured benefits
Forecast failure triggers reassessment	A failed state-gain forecast does not prove avoidance; it lowers confidence in the model and prompts broader inquiry

The last two rules correct a real defect in earlier drafts, and the correction has teeth. An earlier formulation reclassified restoration as avoidance when it repeatedly failed to produce a forecast state gain. That inference is invalid on three grounds. It assumes the measurement is right; it assumes recovery is immediate rather than delayed; and it assumes rest requires a productivity justification, which is false as a matter of rights and not merely as a matter of measurement. People retain claims to sleep, health, leisure, and care that do not depend on downstream output. A failed forecast is evidence about the **model**, not a verdict about the person.

Willingness is permitted to do exactly three things: trigger bounded initiation, because predicted aversion poorly predicts experienced aversion; enter the switching-cost term, because starting an aversive task genuinely costs more; and escalate to the deferral tree after repeated deferral. It may never lower a readiness estimate, raise a demand estimate, or make a task ineligible.

Table 8. Ten reasons a task is not being done.

Reason	Diagnostic signal	Appropriate response
Unreadiness	Declared binding condition not met; bounded initiation provides task-relevant evidence	Defer to next adequate window; decompose
Elevated cost	Binding condition met; effort substantially above baseline	Proceed with reduced scope or shorter block
Aversion	Binding condition met;	Proceed; record a preference,

Reason	Diagnostic signal	Appropriate response
	performance normal in bounded initiation	not a capacity
Low expected value	Value genuinely low on inspection	Re-rank or delete the task
Anxiety	Anticipatory distress; performance intact once started	Bounded initiation; reduce ambiguity
Ambiguity	Cannot state the first concrete action	Decompose until step one is stateable
Skill deficit	Errors persist across states	Training, assistance, or delegation — scheduling will not fix this
Environmental obstruction	Blocked by tools, access, or another party	Remove the obstruction; the task was never eligible
Moral refusal	The person believes the task should not be done	Escalate to purpose review, not to the deferral ladder
Rational refusal	The task is illegitimate, unsafe, or outside the person's authority	Refuse and record; this is a correct output

The last two rows are load-bearing. A system that treats every non-performance as a scheduling failure will grind a person against a task they are right to decline. Appendix D therefore checks refusal **first**.

8. MathGov and RippleLogic as Source Architecture

The qualify-before-rank structure derives from the MathGov / RippleLogic framework (McGaughran, 2026). In this paper the architecture is translated into discipline-neutral scientific terms so that the argument survives a reader who rejects the vocabulary. MathGov supplies an auditable governance grammar; it does **not** supply empirical evidence that STA works.

Table 9. Source architecture and its scientific translation.

Source element	STA translation	Status in the decision
Reality Grounding	Declare what is known, inferred, self-reported, task-tested, and uncertain; maintain calibration and claim boundaries	Precondition on claim authority — not an option-rejecting gate. It can force narrowing, escalation, exploratory marking, or refusal, but it does not by itself reject an option

Source element	STA translation	Status in the decision
Rights Floor	Protect sleep, health, dignity, privacy, agency, accommodation, non-coercion, contestability, before optimization	Option-rejecting, lexicographic
Tail-Risk Constraint	Binding readiness, impact, detectability, reversibility, independent review, refusal-to-finalize	Option-rejecting, lexicographic
Containment and Structural Viability	Burden, adherence, autonomy, coordination, switching, failure handling, adequacy of the option set	Option-rejecting; returns a status, not a bit, and may pass with controls
Ripple ranking	Compare qualified options over time horizons and affected parties without allowing benefit to erase protected constraints	Applies only to survivors

The status column is not decoration. Writing Reality Grounding as a rejection term — as earlier formal drafts did — misstates its function and would cause an implementation to discard options that grounding merely wished to mark as exploratory. Note also that Eligible() in the qualification model correctly omits a grounding term: grounding conditions the *evidence* entering the readiness bound, not the eligibility product.

Self-binding claim. The source architecture explicitly separates architectural conformance from empirical validation and from framework superiority. STA adopts the same discipline: a coherent architecture is not evidence that the routing policy improves outcomes. The research program is separately audited against these gates in the accompanying v1.6 completion report. Self-application demonstrates constraint and consistency, not correctness.

The full apparatus — union scopes, welfare dimensions, ripple scoring, sentience-gradient handling, AI deployment controls, provenance certification — is deferred to companion governance work (the companion governance paper).

9. Confirmatory Program, Exploratory Questions, and Revision Rules

Version 1.6 retains four confirmatory claims. Measurement viability is a prior gate, not another hypothesis: if Phase 2 cannot produce task-family-specific evidence with adequate within-person reliability, criterion validity, calibration, burden, and sensitivity to change, the routing program does not proceed. Every confirmatory study must define design-specific smallest effects of interest, non-inferiority margins, and safety bounds from measurement reliability, baseline risk, burden, consequence, and stakeholder judgment.

9.1 Four confirmatory claims

C1 - Within-person correspondence. On preregistered standardized task families, readiness evidence and person-relative demand jointly improve held-out objective outcome prediction over task-only, person-only, sleep/time-only, and recent-performance baselines by the design-specific margin.

C2 - Control mechanism. A specified control changes its preregistered mechanism - demand, prevention, detectability, reversibility, scope, or qualified delegation - enough to exceed its burden. Effects do not generalize to untested controls or tasks.

C3 - Policy advantage. The complete STA policy exceeds the strongest simple contingent heuristic on a preregistered burden-adjusted policy endpoint. The comparator must receive equal attention, planning support, and implementation-intention structure.

C4 - Safety, autonomy, and equity non-inferiority. STA remains within primary margins for sleep disruption, avoidance, distress, self-monitoring burden, autonomy loss, inequity, relational harm, and critical errors. Any exceeded safety margin overrides performance gains.

9.2 Exploratory mechanism and boundary questions

- Do selected standardized task families show reproducible change-points, soft thresholds, or other non-linear response surfaces after measurement error and person heterogeneity are modeled?
- Do insight tasks reproduce reversed or non-monotonic relations with inhibitory control or circadian optimality?
- When does low-demand maintenance reduce later setup friction, and when does it displace needed recovery?
- Does calibration improve with repeated task-relevant feedback, and does it vary by task family?
- Does age moderate the usefulness of clock/chronotype comparators? The reported 45% versus 83% study proportions are hypothesis-generating, not a direct interaction estimate.
- Does personalization outperform population rules after sufficient data, and where does low task discretion eliminate routing value?
- At what switching cost does dynamic routing become inferior to stable blocking?

9.3 Revision rules

- If the Phase 2 measurement gate fails, STA remains a task-analysis and governance framework rather than a readiness-routing instrument.
- If C1 fails, STA cannot claim a readiness-demand mechanism for that task family.
- If C2 fails, the tested safeguard remains ordinary risk advice rather than a demonstrated mechanism that changes admissibility.
- If C3 fails, the public protocol collapses to the best simple heuristic.
- If C4 fails, deployment pauses or is withdrawn regardless of performance gains.
- If exploratory change-point models fail or are not recoverable, natural-threshold language is removed; consequence-based decision cutoffs may remain if their normative or professional warrant is stated accurately.

10. Staged Research Program

10.1 Phase 0 - reproducible scoping review and claim ledger

Register searches across temporal fit, operator functional state, adaptive function allocation, fatigue and vigilance, recovery, implementation intentions, procrastination, safety controls, algorithmic management, disability and accommodation, and human-AI allocation. Publish databases, dates, strings, screening decisions, exclusions, citation verification, and a claim-to-source ledger. The review must be completed before presenting the literature base as comprehensive.

10.2 Archived design diagnostics - what they do and do not show

Successive exploratory simulations identified genuine design risks: flexible hinge models can win spuriously without an explicit null; supplying the true change-point is optimistic; pooled estimation does not test a within-person multilevel theory; calibration is unstable at small Monte Carlo sample sizes; person-varying cutoffs reduce detection; quantized readiness evidence loses information; and results are dominated by measurement error. One supplied null-family script also used Python's randomized hash() function in seed construction, preventing exact reproducibility across processes. Across four exploratory simulation efforts using different estimators and null families, change-point recovery remained weak at plausible measurement noise, especially with person-varying cutoffs. This is a converging negative finding for the tested diagnostic approaches, not evidence that every possible change-point design is impossible.

These diagnostics justify a subtraction, not a positive claim. They do **not** establish statistical power, a minimum participant count, the existence of a threshold, or the impossibility of all binary-outcome designs. The archived code and outputs are retained as a correction ledger. They are not evidential support for STA.

10.3 Phase 2 - construct and measurement viability gate

Measurement is the decisive phase. Develop task-family-specific readiness and demand measures through cognitive task analysis, participant interviews, expert ratings, behavioral probes, and multilevel psychometrics. Estimate within-person reliability, criterion validity, measurement invariance, calibration, reactivity, missingness, and burden. Quantify the measurement-error level that subsequent mechanism and policy studies would inherit.

Predefine a go/no-go rule. If readiness evidence does not add practically meaningful out-of-sample prediction beyond recent task performance, sleep/time, task identity, and simple self-report—or if obtaining it creates excessive burden or reactivity—the state-routing program pauses. Passive physiological sensing remains a separate high-risk research path; it is not the default response to noisy self-report.

10.4 Phase 3 - descriptive intensive longitudinal study

Phase 3 is for feasibility, measurement structure, calibration, naturalistic selection, and transition description. Because participants choose tasks, it cannot establish the causal value of matching. Use it to refine task families, estimate variance components, quantify missingness and burden, and design controlled studies. Any functional-form analysis is exploratory.

10.5 Phase 4 - controlled correspondence and control-mechanism studies

Use standardized repeated tasks, manipulated or naturally instrumented state variation, blinded outcome scoring, active controls, and analysis matched to the measurement structure. C1 requires continuous or high-information objective outcomes such as detection rate, reconciliation accuracy, response-surface quality, or blind-rated draft quality. C2 requires mechanism-specific experiments that separately test demand reduction, prevention, detectability, reversibility, containment, or delegation. Age and chronotype should be stratification or prespecified moderation variables rather than post hoc explanations.

10.6 Phase 5 - comparative policy and safety trials

Compare full STA, the strongest simple contingent heuristic, priority-only planning, and scheduling-only active controls. Policy outcomes include meaningful completion, objective quality, delay, burden, and sustainable functioning. Critical errors, adverse events, rework, sleep disruption, avoidance, autonomy, distress, inequity, and relational effects are safety outcomes powered and interpreted separately. Safety and equity margins are primary.

10.7 Phase 6 - long-term governance and implementation

Only after evidence from prior phases should organizational or AI-mediated studies proceed. They require worker governance, non-punitive use, accommodation, privacy and security analysis, legal review, opt-out without retaliation, auditability, human override, and monitoring for productivity extraction and structural blame.

11. Measurement and Analysis

11.1 Minimal decision-point scan

The scan should take under thirty seconds and be used **only at meaningful block boundaries**, not continuously. In practical use, only the dimensions binding for the candidate task need rating — typically two or three, not nine. The nine-item version below is the **research** instrument; conflating it with the practical protocol would make STA appear far more burdensome than it needs to be, and would itself be a viability failure.

Table 11. Research scan.

Domain	Prompt	Scale
Integrative clarity	Can I understand, combine, and hold the information	0–4 readiness

Domain	Prompt	Scale
	needed for this block?	
Executive control	Can I begin, resist distraction, and persist for the planned block?	0–4 readiness
Affect regulation	Can I tolerate frustration or uncertainty without becoming reactive?	0–4 readiness
Precision / error monitoring	Am I likely to notice small inconsistencies and check details patiently?	0–4 readiness
Physical readiness	Can I safely perform the physical component?	0–4 readiness
Activation	How activated am I, independent of pleasantness?	0–4 modifier
Associative breadth	How broadly and flexibly are ideas connecting?	0–4 modifier
Willingness	How willing am I to invest effort in <i>this</i> episode?	0–4 separate
Evidence quality	How well is this supported by sleep, context, recent performance, calibration history?	Low / medium / high

Burden is an outcome, not an assumption. Four to six prompts daily over three to four weeks generates roughly 900–1,300 item responses per participant before task ratings. the accompanying audit subjects this to the framework’s own viability gate rather than asserting acceptability.

11.2 Who rates what

The rater assignment is asymmetric, and the asymmetry is the point.

- **Execution risk and reversibility** should be rated by domain experts or by documented rule, **not** by the person under time pressure. A person who wants to finish has an incentive to underrate irreversibility, and this is the most exploitable point in the model.
- **Demand** is participant-rated, calibrated against expert task analysis on a reference subset, and updated from observed performance.
- **Binding sets** are initially expert-assigned by task family, then learned from which shortfalls actually predict failure.

Circularity guard. If demand is inferred from the same outcomes the model predicts, C1 is guaranteed. Demand ratings used in confirmatory tests must be fixed before outcome observation. This is a preregistration requirement, not an analytic preference.

11.3 Analysis strategy

Early analyses should distinguish within-person from between-person effects and must **avoid regressing outcomes on a composite fit score while simultaneously controlling for all components from which that score is deterministically constructed**.

Multilevel models, generalized additive models, Bayesian hierarchical models, dynamic structural equation models, person-specific time-series models, and micro-randomized treatment-effect models are each appropriate to different phases. The central empirical object is not a universal score; it is the estimated **outcome surface** linking person-relative demands, state evidence, controls, and task-specific outcomes.

Exploratory functional-form analyses must test non-linearity **without fishing for cut points**. Decision cutoffs established by professional standards are analyzed separately from empirically estimated change-points. Any proposed natural change-point must be estimated with nested cross-validation and external validation, with measurement error and person heterogeneity modeled explicitly. Personalization is compared against population models on held-out episodes.

12. Applications, Boundaries, and Ethical Governance

12.1 Current defensible scope

The least risky initial scope is voluntary personal experimentation with task decomposition, protected focus windows, source checking, delayed finalization, and explicit non-selection. Even this is unvalidated and must not displace medical, legal, occupational, educational, or professional advice.

12.2 Structural discretion and organizational risk

STA presupposes a meaningful option set and authority over sequencing. Where work is externally paced, algorithmically assigned, understaffed, or deadline-saturated, mismatch is often a job-design problem. Job demands-resources theory provides an established organizational account of how demands and resources shape strain and engagement, and STA must not individualize those structural relations (Bakker & Demerouti, 2017). The correct output may be **DECLARE STRUCTURAL CONSTRAINT**, protected recovery, added staffing, redesigned sequence, mandatory cross-checking, or authority to refuse finalization - not an instruction to self-optimize more efficiently.

Workplace deployment creates an extraction pathology: peak states can be captured for maximal output, low states filled with maintenance, and restoration justified only by later productivity. Rights to sleep, care, disability accommodation, leisure, privacy, and non-instrumental time must remain outside the objective.

12.3 EU legal boundary as of August 2026

This section is a research boundary, not legal advice. Regulation (EU) 2024/1689 defines emotion recognition as biometric inference of emotions or intentions and expressly excludes

physical states such as pain or fatigue. The workplace and education prohibition on biometric emotion inference therefore does not automatically prohibit fatigue detection. The construct boundary matters: biometric inference of affect regulation or reactivity may be closer to prohibited emotion inference than measurement of physical fatigue, and no STA design should infer affect regulation biometrically in workplaces or educational institutions.

Annex III, point 4(b), expressly lists AI systems intended to allocate work tasks on the basis of individual behavior, personal traits, or characteristics, or to monitor and evaluate workers' performance or behavior. An organizational STA system with that intended purpose enters the Annex III classification pathway; whether a particular narrow procedural or preparatory system qualifies for an Article 6 exemption requires a documented intended-purpose and risk analysis. A voluntary personal planning tool that does not influence employment decisions is a materially different use case.

Regulation (EU) 2026/1744, effective 27 July 2026, moved application of the principal Annex III high-risk obligations to 2 December 2027. The same amendment added Article 5 prohibitions concerning non-consensual intimate material and child sexual abuse material from 2 December 2026; those additions do not materially alter STA's worker-routing analysis. A delayed compliance date is not an ethical permission: data protection, labor, equality, collective-bargaining, occupational-safety, and sector-specific duties can apply independently.

The deferral does not make the intervening period legally empty. Where Article 50 applies, transparency duties — including disclosure that a person is interacting directly with an AI system unless that is obvious in the circumstances — apply from 2 August 2026. Article 4's AI-literacy obligation already applied under Chapter I; Regulation (EU) 2026/1744, in force from 27 July 2026, restated it as a duty for providers and deployers to take measures supporting the development of staff AI literacy without guaranteeing a particular competence level. Thus, Annex III deferral must not be read as absence of present transparency, literacy, privacy, labor, equality, safety, or sector-specific obligations (Regulation (EU) 2024/1689; Regulation (EU) 2026/1744).

The design implications are:

1. prefer disclosed voluntary self-report and task evidence over covert inference;
2. exclude biometric affect-regulation inference from workplace and education pathways;
3. treat worker task-allocation or monitoring systems as presumptively within the Annex III pathway unless a documented exemption applies;
4. separate legal classification, scientific validity, and ethical acceptability;
5. require current jurisdiction-specific review before any organizational study or deployment.

12.4 Rights floor

Risk	Minimum safeguard
Coercive monitoring	Genuine opt-in, costless opt-out, no retaliation, worker representation
Sensitive inference	No covert emotion, psychiatric, health, or readiness inference; strict purpose limitation

Risk	Minimum safeguard
Managerial misuse	and minimization No individual state data for hiring, promotion, discipline, insurance, or punitive scheduling
Productivity extraction	Protected sleep, recovery, care, leisure, and non-work time
Disability discrimination	Accommodation and job redesign before individual attribution
Algorithmic opacity	Explanation, uncertainty, contestability, human override, auditable logs
Authority transfer	Recommendation remains separate from authority to sign, submit, prescribe, send, or commit
Structural injustice	Report inadequate staffing, scheduling, tools, or option sets rather than blaming the person

12.5 Safety-sensitive work

In safety-critical settings, STA could at most add a conservative veto or safeguard signal after task-specific validation. It cannot replace fatigue-risk management, duty-time rules, staffing, professional competence, checklists, permits to work, redundancy, supervision, or independent verification. A person's feeling of readiness never waives an institutional safety duty.

12.6 Clinical and disability boundaries

STA is not diagnostic or therapeutic. Persistent changes in capacity should not be absorbed indefinitely as scheduling problems. Research with disabled, neurodivergent, chronically ill, or clinically distressed participants must be co-designed, avoid deficit assumptions, and test whether the instrument creates burden or discrimination. The mechanism and accommodations matter; diagnostic-group labels do not authorize routing decisions.

13. Worked Cases

Each case shows the decision under the current architecture and, where instructive, what a **compensatory** model would have done. The contrast is the argument for the qualification model.

13.1 High-consequence administrative application

A de-identified composite. Any long-form official application with legal effect has this structure.

Component	Binding demand	Risk profile	Routing and control
Collect and rename evidence	Executive control	Low impact; highly detectable; reversible	Lower readiness suffices; checklist;

Component	Binding demand	Risk profile	Routing and control
Transfer verified names and dates	Precision, executive control	Moderate delay risk; reviewable	preserve originals Stable moderate readiness; compare directly against source documents
Interpret eligibility questions	Clarity, precision, regulation	High impact; low detectability	Meet declared binding conditions; use official definitions and expert verification
Reconcile cross-form histories	Precision / error monitoring	High contradiction risk; propagates	High precision; cross-document table; delayed second pass
Final declaration or submission	Precision, regulation	Very high impact; low reversibility	Fresh review window; independent check; refuse finalization under doubtful readiness or low evidence quality

“Paperwork” is not one task class. High readiness may rationally be allocated to tedious protective work because consequence and reversibility dominate present choice. Note that the highest control level attaches to interpretation and submission — the components with the *least* visible work — because detectability and reversibility, not effort, determine the risk profile.

Willingness in this case is low and readiness is adequate. Under Table 7, aversion triggers bounded initiation and enters the switching cost. It does nothing else. The task proceeds.

13.2 Generative work under partial readiness

Clarity and associative breadth are high; precision monitoring is low.

Drafting an argument qualifies. Reference verification and numerical checking do not. The control bundle is to draft with **explicit uncertainty markers** and reserve verification for a later precision episode. A whole-project label would obscure this safe decomposition and would force an all-or-nothing choice between doing the work badly and not doing it.

This case also illustrates why associative breadth is a modifier rather than a capacity: the same high value that supports drafting is a liability for the verification component in the same project.

13.3 Relational action under dysregulation

A difficult conversation may be cognitively clear but emotionally unsafe. Affect regulation is designated binding; the available evidence does not meet the decision condition, and activation is elevated beyond the task-appropriate range.

Cognitive surplus does not authorize the conversation. Eligible alternatives are drafting private notes, gathering facts, scheduling the discussion, or seeking mediation.

What a compensatory model would have done. Under a weighted-sum fit score with an exponential penalty, high importance and high urgency could outrank the regulation deficit, because such a score never reaches zero. The conversation would have been recommended. Under the binding condition in the readiness model it cannot be, at any value — and the activation band supplies a second, independent block. This case is the clearest demonstration that STA is not a productivity maximizer with safety language attached.

13.4 Low readiness: maintenance, restoration, and the forecast trap

After illness or sustained strain, file organization may reduce tomorrow's setup cost — or may consume scarce recovery. The decision depends on expected burden, recovery need, deadline, and intrinsic health value.

Suppose restoration is chosen and the forecast state gain does not materialise. **REVIEW asks whether the forecast was wrong. It does not relabel needed rest as avoidance.** Under an earlier formulation, repeated forecast failure would have reclassified restoration as avoidance and escalated. That inference assumed the measurement was right, that recovery is immediate, and that rest requires a productivity justification. All three are questionable, and the third is a rights claim rather than an empirical one. The correct response is reduced confidence in the model and broader inquiry — including, per Appendix D, a health check where low readiness persists.

14. Discussion

STA began as practical advice about matching high and low energy to different work. The mature question is narrower and more demanding: can a rights-constrained policy select task-control bundles better than simple contingent planning when readiness is uncertain, task demands are person-relative, and consequences are asymmetric?

Version 1.6 makes three separations that the earlier manuscript lacked. First, descriptive science is not the routing policy: a readiness-outcome relation does not prove that a complex scheduler improves life. Second, the routing policy is not the rights floor: empirical benefit cannot authorize coercive monitoring, sleep extraction, or the loss of protected care and non-work. Third, mechanism outcomes are not safety outcomes: continuous repeated quality measures may test a change-point, while rare critical errors test whether deployment is safe.

The strongest surviving contribution is therefore not a universal threshold model. It is a disciplined integration: task-episode decomposition, uncertain evidence, person-relative demand, mechanism-specific safeguards, qualify-before-rank ordering, explicit non-selection,

and comparative tests against simpler policies. Each component has antecedents. The work earns value only if the integration produces transparent, reproducible, and rights-respecting policy advantages.

The greatest vulnerability remains measurement. Brief self-reports may reflect a general feeling-capable factor, task demand is endogenous to skill and tools, and the person selects both task and timing. Physiological sensing does not automatically solve this problem; the operator-functional-state literature is heterogeneous and more successful at detecting state correlates than predicting operational performance. The appropriate response is task-specific measurement, controlled repeated tasks, transparent uncertainty, and conservative claims.

The simulation program is a credibility gain because discouraging diagnostics changed the theory. It is not evidence for STA and is not a completed power analysis. Across iterations, results proved sensitive to measurement error, null-family choice, calibration stability, seed construction, and person-varying cutoffs. The correct response is the subtraction made here: sharp functional form is exploratory, measurement is the go/no-go gate, and no simulation-derived participant count is presented as established.

A final risk is totalizing optimization. A system that turns every state into an input for productive routing can colonize the life it claims to help. Care/non-work, refusal, protected rest, and structural redesign are not residual categories. They are proof that the framework governs optimization rather than worships it.

15. Conclusion and Publication Path

Human readiness varies, tasks impose different person-relative demands, safeguards alter risk, and self-knowledge is imperfect. These premises justify a research problem. They do not yet validate a scheduling instrument.

STA now consists of three linked but independently accountable layers. The descriptive layer asks which readiness-demand-control relations predict objective outcomes. The prescriptive layer asks whether routing task-control bundles improves meaningful and sustainable performance beyond simple contingent rules. The normative layer refuses to trade protected rights, severe risk, care, or structural justice for aggregate productivity.

The strongest empirical program is now clear:

1. complete and publish a reproducible scoping review and verification ledger;
2. make task-family-specific measurement a formal go/no-go gate rather than assuming readiness can be measured adequately;
3. test readiness-demand correspondence and control mechanisms on standardized objective outcomes;
4. compare the full policy against the strongest simple contingent heuristic;
5. power safety, autonomy, equity, and burden outcomes separately and treat them as primary constraints;
6. simplify or withdraw the public protocol when the evidence requires it;

7. keep organizational and AI deployment in a separate governance program until evidence and law support it.

The central test is not whether the architecture sounds wise. It is whether it produces better decisions, under honest uncertainty, without converting human variability into a new instrument of surveillance, extraction, avoidance, or blame. Until that test is passed, STA is a disciplined research architecture - not a claim to have solved human energy management, discovered a universal threshold, or validated a readiness score.

Appendix A. Minimal State and Evidence Card

Readiness / modifier	0	2	4	Evidence anchor
Integrative clarity	Cannot hold the problem	Adequate for familiar work	Can integrate complex material	Recent comprehension or reasoning behaviour
Executive control	Cannot initiate or persist	Can sustain a bounded routine block	Strong initiation and persistence	Recent starts, switches, completions
Affect regulation	Highly reactive or dysregulated	Can tolerate ordinary frustration	Stable under uncertainty or conflict	Recent behaviour, not intended composure
Precision monitoring	Likely to miss obvious errors	Adequate with checklist	Sustained careful checking	Recent error detection or review sample
Physical readiness	Unsafe or unable	Adequate for light component	Strong and safe	Symptoms, exertion, context
Activation	Very low	Moderate	Very high	<i>Modifier.</i> Optimal level is task-dependent
Associative breadth	Narrow or rigid	Ordinary	Broad or flexible	<i>Modifier.</i> May aid insight and impair precision
Willingness	Strong refusal or aversion	Mixed	High willingness	Separate from ability. Never adjusts a readiness rating

Evidence quality: low / medium / high. Record sleep and context, recent task-relevant performance, known calibration error, and material uncertainty. Low evidence quality does not automatically prohibit work; it raises safeguards for consequential and hard-to-reverse episodes and can make finalization ineligible.

Appendix B. Task-Control Card

- **Component and scope.** What exactly will be done in this block? What will *not* be finalized?
- **Binding demands.** Which readiness dimensions or professional conditions are designated non-compensatory for this episode, and why?
- **Execution risk.** What errors can occur, who is affected, how detectable, how reversible, how far can harm propagate?
- **Value.** What benefit is created and what omission loss avoided, over which horizon, for which parties?
- **Constraints.** Deadline, minimum block, setup, switching, sequence, coordination, autonomy.
- **Controls.** Checklist, source trace, second reader, delay, staged action, reduced scope, refusal to finalize.
- **Decision.** Select · select with controls · prepare only · defer to next adequate window · restore · care / non-work · escalate · declare structural constraint.

Appendix C. Adoption Gates

Use level	Minimum evidence before recommendation
Personal experimental use	Feasibility, burden, basic calibration, transparent non-validated framing, no displacement of health care or professional advice
Education	Construct validity, comparative trials, age-appropriate protections, teacher workload analysis, adverse-effect monitoring, non-punitive policy, no affect inference
Organizations	Comparative field evidence, enforceable rights floor, independent worker governance, privacy safeguards, accommodations, non-punitive use, legal review
Safety-sensitive work	Task-specific validated risk limits or professional standards, institutional fatigue systems, independent controls, and authority boundaries. Veto signal only; not a routing input at current evidence levels
AI-mediated routing	All applicable evidence plus fairness, automation-bias, security, privacy, override, audit, and no autonomous consequential authority . Biometric emotion inference at work or in education is prohibited in the EU; fatigue and other physical-state systems require separate classification and legal analysis

Appendix D. Deferral Decision Tree

Triggered on the second deferral of any task with high value or high omission loss.

Task deferred twice.

- 1. Is refusal correct?
Illegitimate, unsafe, outside my authority, or something I believe should not be done?
 - └ YES → REFUSE. Record the reason. Escalate to purpose review.
Do NOT proceed down this tree. The deferral was correct.
- 2. Is the obstruction external?
Blocked by access, tools, information, or another party?
 - └ YES → The task was never eligible. Remove the obstruction or reassign. This is not a deferral.
- 3. Is it specifiable?
Can I state the first concrete action in one sentence?
 - └ NO → AMBIGUITY. Decompose until step one is stateable.
Re-enter routing with the sub-component.
- 4. Run bounded initiation (10–15 min, reversible).
Never overrides a rights or safety gate.
 - └ Performance normal → AVERSION, not unreadiness.
Continue to the block boundary.
Record a preference, not a capacity.
 - └ Errors or non-comprehension → UNREADINESS.
Reserve the next adequate window.
Record the reservation.
 - └ Errors persist across states → SKILL GAP.
Training, assistance, delegation.
Scheduling will not fix this.
- 5. Deferred three or more times despite meeting the declared conditions once?
 - └ YES → ESCALATE. Choose one explicitly and record it:
 - (a) decompose further
 - (b) obtain assistance or delegate
 - (c) change the environment
 - (d) renegotiate deadline or scope
 - (e) accept the cost of not doing it, in writingSilent continued deferral is not an option.

- └ 6. Health check — unconditional.
Has readiness evidence remained materially impaired for most windows across two weeks or more?
└ YES → This is not a scheduling problem.
Seek medical, psychological, or occupational assessment.
STA does not diagnose and must not be used to manage a persistent capacity change.

Two design features are deliberate. **Refusal is checked first**, so the escalation ladder can never be applied to a task the person is right to decline. **The health check is unconditional and last**, so a pattern of low readiness cannot be indefinitely absorbed as an allocation problem.

References

- Albulescu, P., Macsinga, I., Rusu, A., Sulea, C., Bodnaru, A., & Tulbure, B. T. (2022). “Give me a break!” A systematic review and meta-analysis on the efficacy of micro-breaks for increasing well-being and performance. *PLOS ONE*, 17(8), e0272460. <https://doi.org/10.1371/journal.pone.0272460>
- Baas, M., De Dreu, C. K. W., & Nijstad, B. A. (2008). A meta-analysis of 25 years of mood-creativity research: Hedonic tone, activation, or regulatory focus? *Psychological Bulletin*, 134(6), 779–806. <https://doi.org/10.1037/a0012815>
- Bakker, A. B., & Demerouti, E. (2017). Job demands-resources theory: Taking stock and looking forward. *Journal of Occupational Health Psychology*, 22(3), 273–285. <https://doi.org/10.1037/ocp0000056>
- Bennett, A. A., Gabriel, A. S., & Calderwood, C. (2020). Examining the interplay of micro-break durations and activities for employee recovery: A mixed-methods investigation. *Journal of Occupational Health Psychology*, 25(2), 126–142. <https://doi.org/10.1037/ocp0000168>
- Blain, B., Hollard, G., & Pessiglione, M. (2016). Neural mechanisms underlying the impact of daylong cognitive work on economic decisions. *Proceedings of the National Academy of Sciences*, 113(25), 6967–6972. <https://doi.org/10.1073/pnas.1520527113>
- Boksem, M. A. S., & Tops, M. (2008). Mental fatigue: Costs and benefits. *Brain Research Reviews*, 59(1), 125–139. <https://doi.org/10.1016/j.brainresrev.2008.07.001>
- Cao, Y., Xie, T., & Ma, N. (2025). The impairments of sleep loss on core executive functions: General and task-specific effects. *Sleep Medicine Reviews*, 84, 102163. <https://doi.org/10.1016/j.smr.2025.102163>
- Chauhan, S., Vanova, M., Tailor, U., Asad, M., Faßbender, K., Norbury, R., Ettinger, U., & Kumari, V. (2025). Chronotype and synchrony effects in human cognitive performance: A systematic review. *Chronobiology International*, 42(4), 463–499. <https://doi.org/10.1080/07420528.2025.2490495>

Cohen, A.-L., Bayer, U. C., Jaudas, A., & Gollwitzer, P. M. (2008). Self-regulatory strategy and executive control: Implementation intentions modulate task switching and Simon task performance. *Psychological Research*, 72(1), 12–26. <https://doi.org/10.1007/s00426-006-0074-2>

Dust, S. B. (2015). Mindfulness, flow, and mind wandering: The role of trait-based mindfulness in state–task alignment. *Industrial and Organizational Psychology*, 8(4), 609–614. <https://doi.org/10.1017/iop.2015.87>

Edwards, J. R., & Parry, M. E. (1993). On the use of polynomial regression equations as an alternative to difference scores in organizational research. *Academy of Management Journal*, 36(6), 1577–1613. <https://doi.org/10.2307/256822>

George, J. M., & Zhou, J. (2002). Understanding when bad moods foster creativity and good ones don't: The role of context and clarity of feelings. *Journal of Applied Psychology*, 87(4), 687–697. <https://doi.org/10.1037/0021-9010.87.4.687>

Glöckner, A. (2016). The irrational hungry judge effect revisited: Simulations reveal that the magnitude of the effect is overestimated. *Judgment and Decision Making*, 11(6), 601–610. <https://doi.org/10.1017/S1930297500004812>

Gollwitzer, P. M., & Sheeran, P. (2006). Implementation intentions and goal achievement: A meta-analysis of effects and processes. *Advances in Experimental Social Psychology*, 38, 69–119. [https://doi.org/10.1016/S0065-2601\(06\)38002-1](https://doi.org/10.1016/S0065-2601(06)38002-1)

Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., ... Zwienerberg, M. (2016). A multilab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, 11(4), 546–573. <https://doi.org/10.1177/1745691616652873>

Health and Safety Executive. (n.d.). Permit to work systems. <https://www.hse.gov.uk/humanfactors/topics/ptw.htm>

Hockey, G. R. J. (2013). *The psychology of fatigue: Work, effort and control*. Cambridge University Press.

Inzlicht, M., & Schmeichel, B. J. (2012). What is ego depletion? Toward a mechanistic revision of the resource model of self-control. *Perspectives on Psychological Science*, 7(5), 450–463. <https://doi.org/10.1177/1745691612454134>

Jurgelis, M., Chong, W. B., Atkins, K. J., Cooper, P. S., Coxon, J. P., & Chong, T. T.-J. (2021). Heightened effort discounting is a common feature of both apathy and fatigue. *Scientific Reports*, 11, 22283. <https://doi.org/10.1038/s41598-021-01287-2>

Kaber, D. B., & Endsley, M. R. (2004). The effects of level of automation and adaptive automation on human performance, situation awareness and workload in a dynamic control task. *Theoretical Issues in Ergonomics Science*, 5(2), 113–153. <https://doi.org/10.1080/1463922021000054335>

Klasnja, P., Hekler, E. B., Shiffman, S., Boruvka, A., Almirall, D., Tewari, A., & Murphy, S. A. (2015). Microrandomized trials: An experimental design for developing just-in-time adaptive interventions. *Health Psychology*, 34(Suppl.), 1220–1228. <https://doi.org/10.1037/hea0000305>

- Kool, W., & Botvinick, M. (2014). A labor/leisure tradeoff in cognitive control. *Journal of Experimental Psychology: General*, 143(1), 131–141. <https://doi.org/10.1037/a0031048>
- Kristof-Brown, A. L., Zimmerman, R. D., & Johnson, E. C. (2005). Consequences of individuals' fit at work: A meta-analysis of person–job, person–organization, person–group, and person–supervisor fit. *Personnel Psychology*, 58(2), 281–342. <https://doi.org/10.1111/j.1744-6570.2005.00672.x>
- Kurzban, R., Duckworth, A., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, 36(6), 661–679. <https://doi.org/10.1017/S0140525X12003196>
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43, e1. <https://doi.org/10.1017/S0140525X1900061X>
- Lowe, C. J., Safati, A., & Hall, P. A. (2017). The neurocognitive consequences of sleep restriction: A meta-analytic review. *Neuroscience & Biobehavioral Reviews*, 80, 586–604. <https://doi.org/10.1016/j.neubiorev.2017.07.010>
- Mancuso, A., Compare, M., Salo, A., & Zio, E. (2017). Portfolio optimization of safety measures for reducing risks in nuclear systems. *Reliability Engineering & System Safety*, 167, 20–29. <https://doi.org/10.1016/j.res.2017.05.005>
- May, C. P., Hasher, L., & Stoltzfus, E. R. (1998). Synchrony effects in inhibitory control over thought and action. *Journal of Experimental Psychology: Human Perception and Performance*, 24(2), 363–379. <https://doi.org/10.1037/0096-1523.24.2.363>
- McCauley, M. E., McCauley, P., Riedy, S. M., Banks, S., Ecker, A. J., Kalachev, L. V., Rangan, S., Dinges, D. F., & Van Dongen, H. P. A. (2021). Fatigue risk management based on self-reported fatigue: Expanding a biomathematical model of fatigue-related performance deficits to also predict subjective sleepiness. *Transportation Research Part F: Traffic Psychology and Behaviour*, 79, 94–106. <https://doi.org/10.1016/j.trf.2021.04.006>
- McGaughran, J. (2026). *MathGov Core Release 2026.09, v12.5 / SGP v8.4 — Configuration-Bound Assurance Public Research Source Release* [Open research-source specification]. GitHub. <https://github.com/MathGov/ripple-logic/releases/tag/v12.5>
- Meijman, T. F., & Mulder, G. (1998). Psychological aspects of workload. In P. J. D. Drenth, H. Thierry, & C. J. de Wolff (Eds.), *Handbook of work and organizational psychology: Work psychology* (2nd ed., pp. 5–33). Psychology Press.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7(3), 134–140. [https://doi.org/10.1016/S1364-6613\(03\)00028-7](https://doi.org/10.1016/S1364-6613(03)00028-7)
- Nahum-Shani, I., Smith, S. N., Spring, B. J., Collins, L. M., Witkiewitz, K., Tewari, A., & Murphy, S. A. (2018). Just-in-time adaptive interventions (JITAIs) in mobile health: Key components and design principles for ongoing health behavior support. *Annals of Behavioral Medicine*, 52(6), 446–462. <https://doi.org/10.1007/s12160-016-9830-8>

- Parasuraman, R., Mouloua, M., & Molloy, R. (1996). Effects of adaptive task allocation on monitoring of automated systems. *Human Factors*, 38(4), 665–679. <https://doi.org/10.1518/001872096778827279>
- Pelechano, V., Mestre, A., Albert, M., & Gil, M. (2026). *HAAS: A policy-aware framework for adaptive task allocation between humans and artificial intelligence systems* [Preprint]. arXiv:2605.02832. <https://doi.org/10.48550/arXiv.2605.02832>
- Pütz, S., Mertens, A., Chuang, L. L., & Nitsch, V. (2024). Physiological measures of operators' mental state in supervisory process control tasks: A scoping review. *Ergonomics*, 67(6), 801–830. <https://doi.org/10.1080/00140139.2023.2289858>
- Pütz, S., Mertens, A., Chuang, L. L., & Nitsch, V. (2025). Physiological predictors of operator performance: The role of mental effort and its link to task performance. *Human Factors*, 67(6), 595–615. <https://doi.org/10.1177/00187208241296830>
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence. Official Journal of the European Union, L, 2024/1689. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Regulation (EU) 2026/1744 of the European Parliament and of the Council of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence. *Official Journal of the European Union*, L 2026/1744.
- Rubinstein, J. S., Meyer, D. E., & Evans, J. E. (2001). Executive control of cognitive processes in task switching. *Journal of Experimental Psychology: Human Perception and Performance*, 27(4), 763–797. <https://doi.org/10.1037/0096-1523.27.4.763>
- Sheeran, P., Listrom, O., & Gollwitzer, P. M. (2025). The when and how of planning: Meta-analysis of the scope and components of implementation intentions in 642 tests. *European Review of Social Psychology*, 36(1), 162–194. <https://doi.org/10.1080/10463283.2024.2334563>
- Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, 79(2), 217–240. <https://doi.org/10.1016/j.neuron.2013.07.007>
- Sonnentag, S., & Fritz, C. (2007). The Recovery Experience Questionnaire: Development and validation of a measure for assessing recuperation and unwinding from work. *Journal of Occupational Health Psychology*, 12(3), 204–221. <https://doi.org/10.1037/1076-8998.12.3.204>
- Soutschek, A., & Tobler, P. N. (2020). Causal role of lateral prefrontal cortex in mental effort and fatigue. *Human Brain Mapping*, 41(16), 4630–4640. <https://doi.org/10.1002/hbm.25146>
- Steel, P. (2007). The nature of procrastination: A meta-analytic and theoretical review of quintessential self-regulatory failure. *Psychological Bulletin*, 133(1), 65–94. <https://doi.org/10.1037/0033-2909.133.1.65>
- Steel, P., & König, C. J. (2006). Integrating theories of motivation. *Academy of Management Review*, 31(4), 889–913. <https://doi.org/10.5465/amr.2006.22527462>

Tseng, H. K., & Damian, M. F. (2023). Exploring synchrony effects in performance on tasks involving cognitive inhibition: An online study of young adults. *Chronobiology International*, 40(9), 1209–1223. <https://doi.org/10.1080/07420528.2023.2256843>

Van Dongen, H. P. A., Maislin, G., Mullington, J. M., & Dinges, D. F. (2003). The cumulative cost of additional wakefulness: Dose-response effects on neurobehavioral functions and sleep physiology from chronic sleep restriction and total sleep deprivation. *Sleep*, 26(2), 117–126. <https://doi.org/10.1093/sleep/26.2.117>

Vleugels, W., & Flatau-Harrison, H. (2023). Curvilinear relationships in person–environment fit research: Is there evidence for a too-much-of-a-good-thing effect? *Journal of Occupational and Organizational Psychology*, 96(3), 669–677. <https://doi.org/10.1111/joop.12432>

Vleugels, W., Verbruggen, M., De Cooman, R., & Billsberry, J. (2023). A systematic review of temporal person–environment fit research: Trends, developments, obstacles, and opportunities for future research. *Journal of Organizational Behavior*, 44(2), 376–398. <https://doi.org/10.1002/job.2607>

Vohs, K. D., Schmeichel, B. J., Lohmann, S., Gronau, Q. F., Finley, A. J., Ainsworth, S. E., ... Albarracín, D. (2021). A multisite preregistered paradigmatic test of the ego-depletion effect. *Psychological Science*, 32(10), 1566–1581. <https://doi.org/10.1177/0956797621989733>

Warm, J. S., Parasuraman, R., & Matthews, G. (2008). Vigilance requires hard mental work and is stressful. *Human Factors*, 50(3), 433–441. <https://doi.org/10.1518/001872008X312152>

Wickens, C. D. (2002). Multiple resources and performance prediction. *Theoretical Issues in Ergonomics Science*, 3(2), 159–177. <https://doi.org/10.1080/14639220210123806>

Wickens, C. D. (2008). Multiple resources and mental workload. *Human Factors*, 50(3), 449–455. <https://doi.org/10.1518/001872008X288394>

Wieth, M. B., & Zacks, R. T. (2011). Time of day effects on problem solving: When the non-optimal is optimal. *Thinking & Reasoning*, 17(4), 387–401. <https://doi.org/10.1080/13546783.2011.625663>

Zhang, J.-H., Xia, J.-J., Garibaldi, J. M., Groumpos, P. P., & Wang, R.-B. (2017). Modeling and control of operator functional state in a unified framework of fuzzy inference Petri nets. *Computer Methods and Programs in Biomedicine*, 144, 147–163. <https://doi.org/10.1016/j.cmpb.2017.03.024>

Zijlstra, F. R. H., Cropley, M., & Rydstedt, L. W. (2014). From recovery to regulation: An attempt to reconceptualize recovery from work. *Stress and Health*, 30(3), 244–252. <https://doi.org/10.1002/smi.2604>