



CDERP

CENTRE FOR DISABILITY
EMPLOYMENT RESEARCH
AND PRACTICE

RESEARCH
INSIGHT
PRACTICE
BETTER OUTCOMES

Governing the Moving Maze

Bounded Chaos, Dynamic Equilibrium
and the Limits of System Redesign

PEOPLE
SYSTEMS
OPPORTUNITIES

CDERP Position Paper | October 2026

Melbourne, Australia

Contents

Executive summary	1
1. The attraction of the blueprint	3
2. The system moves while we redesign it.....	5
3. The problem with snapshots	6
4. Dynamic equilibrium.....	8
5. Bounded chaos	10
6. The Fishbowl Effect.....	13
7. When necessary functions escape the architecture	16
8. From events to trajectories	20
9. Feedback as system infrastructure.....	24
10. Governing the moving maze	30
11. An invitation rather than a blueprint	36
Statement on the Use of Artificial Intelligence.....	37
References	37

Suggested citation

Smith, P. (2026). Governing the moving maze: Bounded chaos, dynamic equilibrium and the limits of system redesign. Centre for Disability Employment Research and Practice (CDERP).

Executive summary

Employment, workforce and human-service systems are often redesigned through new arrangements of programs, contracts, incentives and performance measures, as though these could specify how the system will operate in the future. Such architecture is necessary for public accountability, but it cannot determine how people and organisations will respond. Reform unfolds while the system continues to change, and the intervention itself alters some of the conditions it was intended to govern.

The argument developed here is that effective reform requires governance capable of responding to a system in motion, as well as an initial design. Bounded chaos describes movement arising among actors whose authority and responsibilities are bounded. Although institutional boundaries shape that movement, its consequences frequently extend beyond them. Governance cannot produce or regulate movement itself; it can, however, preserve purpose, rights and accountability while allowing necessary information, relationships and corrective responses to cross appropriate boundaries.

Dynamic equilibrium helps explain why recognisable outcomes may persist even as the actors producing them continually adapt. Such stability should not automatically be mistaken for success, since workarounds and informal compensation can sustain an underperforming system while concealing the conditions that reproduce its results. The Fishbowl Effect describes a related difficulty: organisations and programs may perform well within their administrative boundaries even as a person's longer employment trajectory crosses boundaries where knowledge, support and relationships are lost.

An important distinction throughout the paper is between activities and the functions they perform. When redesign removes an activity, the underlying function may still be needed, with the work shifting to practitioners, families, participants or other organisations beyond formal recognition or funding. Functional governance therefore considers which functions remain necessary, where responsibility for them lies and how effectively they are performed, without requiring identical institutional arrangements or assuming that the inventory of functions is fixed.

Administrative events and longer trajectories also require different forms of attention. Job commencement, training completion and retention milestones provide useful observations, but they cannot by themselves establish what an intervention changed over time. Consequential observations must be able to reach actors who can interpret and respond to them. Information collection is not necessarily learning, and even a response intended as an improvement needs to be examined for its subsequent effects.

Rather than proposing a single ideal design, the paper identifies continuing responsibilities of governance: sustaining clarity of purpose, revisiting essential functions, managing boundary permeability, observing trajectories without assigning unlimited accountability, and creating pathways for consequential feedback and proportionate adaptation. The observational models informing these propositions must remain open to revision as further evidence and experience emerge.

The task is not to stop the maze moving. It is to govern it while it moves.

1. The attraction of the blueprint

System redesign usually begins with a legitimate concern that existing arrangements are no longer achieving what is expected of them. Outcomes may be inadequate, expenditure may be rising, people may be falling through gaps, or services may be fragmented. Employers can be dissatisfied, and established arrangements may no longer reflect changing economic, technological or social conditions.

The natural response is to understand the current system and design something better. This typically means identifying deficiencies, defining the outcomes sought, redesigning structures and responsibilities, setting eligibility and funding arrangements, specifying provider obligations, building performance measures and creating mechanisms to align behaviour with the objectives of reform.

This approach has a necessary place in public administration. Complex systems require structure, public resources require accountability, and organisations need to understand their responsibilities. Participants, too, need to know what they can expect, while governments have sound reasons to define the purposes and boundaries of the systems they fund or regulate.

The problem begins when the resulting architecture is treated as though it were the system itself. A program model can specify what should happen. A contract can allocate responsibilities, a performance framework can define what will be measured and a funding model can create incentives. An organisational chart can show formal relationships. None of these can fully determine how the actors within the system will respond to one another over time.^{1,2}

We describe as the blueprint assumption the expectation that a sufficiently detailed map of a system will allow its future state to be specified with comparable precision. It is an attractive assumption because it makes complexity appear administratively manageable: the present can be represented, the desired future designed, and the difference between them treated as the work of reform.

Architecture sets the conditions within which behaviour occurs, but it does not determine that behaviour. Actors interpret rules, respond to incentives, form relationships, protect their interests, exercise judgement and adapt to changing circumstances. The same formal architecture can therefore produce quite different patterns of behaviour across places, organisations and periods.²

Time makes that assumption more difficult to sustain. Reform develops through consultation, drafting, negotiation and implementation, often over considerable periods. As this happens, organisations anticipate new requirements and adjust their behaviour, while technologies, labour markets and wider circumstances continue to change. Some responses to a proposed reform begin well before its formal commencement.

Consequently, the architecture eventually enters a system that has already changed from the one observed at the outset of the redesign.

The blueprint nevertheless remains useful for clarifying intentions, allocating authority and establishing boundaries. Its limitation is that the system cannot be held still while those intentions are translated into practice.

2. The system moves while we redesign it

Employment and workforce systems comprise people and institutions whose actions continually affect one another. Government agencies, providers, practitioners, employers, individuals, families, communities and other services bring different information, responsibilities and incentives to their interactions.

A change affecting one actor changes the circumstances others encounter. A new payment mechanism can shift organisational priorities, which changes practitioner behaviour, which in turn changes participant experience and employer engagement. A change to eligibility can push demand into another service system. A new performance measure can redirect attention towards whatever will be counted. Greater organisational flexibility can generate innovation, and it can also produce inconsistency. Stronger prescription can improve standardisation while constraining adaptation to circumstances nobody anticipated when the rules were written.^{1,3}

Intervention, in other words, does not simply act upon a system; it becomes part of the conditions to which the system responds.

This matters because responses to interventions are not necessarily signs of poor implementation. Poor implementation is real, and some divergence from intended practice does reflect inadequate capability, weak management or non-compliance. Not every divergence can be explained that way, however. Some responses arise because actors are adapting to the new conditions the intervention has itself created.

For example, removing a formal activity does not necessarily remove the need for the function it performed. Where that function remains necessary, practitioners, participants, families or other organisations may take it up elsewhere, so that the redesign changes the location of the work without eliminating it.

Meanwhile, the environment continues to move. Economic conditions alter, employers reorganise, technologies reshape occupations and service delivery, and political priorities change. The reform therefore enters a setting that is already changing, where actors are also responding to the reform as it approaches.

The system receiving the reform will differ from the one initially observed and will continue to change after implementation begins.

This does not make design futile, but it broadens what good design needs to accomplish. Alongside the quality of the initial architecture, we need to consider the capacity to observe consequences, learn from experience and adapt after implementation.

3. The problem with snapshots

Indicators are indispensable to governments and organisations seeking to understand system performance. Employment and unemployment rates, vacancies, commencements, placements, retention, expenditure, caseloads and outcome rates all provide valuable observations, without which much of the system would remain difficult to see.

The difficulty arises when an observation of a system's state is treated as an explanation of the mechanisms producing that state.⁴

An unemployment rate describes measured unemployment, while a vacancy rate describes reported vacancies; neither necessarily explains why unemployed people and vacant jobs coexist. Workers and jobs are not interchangeable units. Occupation, skills, geography, transport, hours, pay, accessibility, recruitment practice, caring responsibilities, working conditions and timing all influence whether a particular person and opportunity can form a viable employment relationship.

System-state statistics describe conditions. They do not, without further evidence, establish the mechanisms that produce those conditions.

The same caution applies in the other direction. Individual experience can reveal mechanisms that aggregate data cannot see, but one person's experience does not establish the condition of the wider system. Different forms of observation answer different questions, and it helps to keep four of them distinct (Figure 1).⁴




Figure 1. *Four forms of observation: system state, mechanism, individual experience and trajectory*

These four perspectives illuminate different aspects of the same system. A system state identifies an observed condition, while a mechanism concerns the processes, relationships and incentives through which that condition is produced or sustained. Individual experience shows how those conditions are encountered, and a trajectory introduces time by asking what happens next.

The distinction matters because an event that appears successful at one point may contribute little to a later trajectory, while an intervention that looks modest in immediate performance data may substantially change what becomes possible later. A snapshot can show that something happened. It cannot, on its own, establish what that event changed.

There is a further complication. Measurement is not always external to the system being measured. Once an indicator becomes consequential, through payment, targets, rankings or accountability, actors respond to it. Attention moves towards what is counted, and activity is reorganised around thresholds and definitions. The measure itself becomes a force within the system.⁵⁻⁷

Intervention becomes part of the system. Measurement can do the same.

Taken together, system states, mechanisms, individual experiences and trajectories offer a fuller understanding than any one perspective can provide. A system may be highly observable through its aggregate indicators while remaining poorly understood in terms of how conditions arise, how people experience them or how their consequences develop over time.

This raises a further question: if the system is continually changing, how should we understand the patterns that nevertheless persist? Dynamic equilibrium offers a way of approaching that problem.

4. Dynamic equilibrium

The Equilibrium Systems Model of Employment (ESME) was developed in response to exactly this problem, drawing on our doctoral systems research and subsequently articulated in CDERP's peer-reviewed ESME paper.^{8,9} Its starting point is that employment outcomes are not produced by a single intervention or actor, but through relationships across multiple domains.

The domains identified in ESME should be understood as an observational architecture rather than as the system itself. They allow interactions among actors, resources, rules, relationships and constraints to be considered together, making connections visible that might otherwise be difficult to examine.

Every model simplifies. Its value lies not in reproducing the system in miniature but in improving the questions we can ask of it. A change within one domain can alter the conditions encountered in others, and the significance of any particular factor can shift as those relationships shift.^{1,3}

An observational framework must also remain open to evidence that challenges its categories and assumptions. Its purpose is to improve our understanding of the system, not to require later observations to conform to the model. As understanding develops, the framework may itself need to be revised, extended or replaced. Analytical continuity should not become theoretical dependence.

Equilibrium in this sense does not mean rest. A person can remain employed while the configuration sustaining that employment changes beneath them. Supervisors move on, support reduces or increases, tasks evolve, family circumstances alter, technologies shift and workplace relationships develop. The observable state persists while the relationships producing it continue to move.^{1,3}

Similar movement occurs across the wider system as providers enter or leave, funding arrangements change, practitioners adapt, employers respond and governments revise policy. Economic conditions also shift, yet familiar patterns of outcomes can persist through these changes.

We use dynamic equilibrium to describe the continuing production and reproduction of an observable condition through interaction, rather than its persistence through an absence of change.

There is no reason to assume that such an equilibrium is desirable. A system may repeatedly produce modest outcomes even while its actors work hard to keep it functioning: organisations adjust to incentives, practitioners navigate gaps, families compensate, functions migrate and governments modify settings. These efforts can sustain the system while leaving the underlying pattern of underperformance largely intact.

Stable underperformance may therefore represent an equilibrium in its own right, rather than a temporary departure from a better state. This matters for reform because repairing an apparently failing component may leave undisturbed the relationships that continue to reproduce the outcome.

It helps to explain why successive reforms can change program names, contracts, organisational structures and performance measures while familiar patterns of outcomes remain. Sometimes the more useful question is not "Which component is failing?" but "What configuration of relationships allows this outcome to be continually reproduced?"

Nor is adaptation necessarily progressive. Workarounds can make poor processes tolerable, families may compensate for missing continuity, and organisations can become adept at meeting performance measures without changing the trajectories those measures were intended to represent. Such compensation can make a system appear more functional than it is, reducing the pressure for structural change.

What appears to be resilience may therefore depend on unrecognised compensation. In some circumstances, the very behaviours that keep an underperforming system operating also contribute to the durability of its underperformance.

Dynamic equilibrium is intended here as a descriptive rather than a normative concept. It draws attention to the relationships through which observable conditions are continually recreated, allowing us to understand apparent stability as something produced through movement.

The governance problem, then, is not simply how to move a system from one static equilibrium to another. It is how to understand the relationships producing its current condition, observe what happens when those relationships are disturbed, and retain the capacity to respond as the system changes.

5. Bounded chaos

Dynamic equilibrium presents a particular governance difficulty. If advance specification cannot fully control a system, it may be tempting to favour maximum flexibility and allow actors to adapt as they see fit. Yet complex institutional systems also depend on structures that make action and accountability possible.

The issue is therefore how structure and movement can coexist without either being mistaken for the whole of governance.

We use the term bounded chaos to describe a characteristic of complex institutional systems: movement continually arises through interactions among actors whose authority, resources and responsibilities are bounded. Chaos here does not mean randomness, disorder or the absence of governance, nor does it imply a mathematically chaotic state. The boundaries shape but do not determine the movement, and governance must accommodate rather than manufacture or maintain it. This account draws on complexity scholarship concerning emergence, the limits of prediction and system representation, and resilience.^{10–12}

Those boundaries take many forms. Legislation, funding, contracts, eligibility rules, professional roles, organisational routines, markets, institutional cultures and social expectations all shape what actors can do. Boundaries distribute authority, constrain behaviour and make coordination possible.^{2,13}

Even so, boundaries cannot determine every interaction. Practitioners exercise judgement, employers interpret risk and opportunity, participants respond to experience, and organisations develop routines. Families may compensate for gaps, information travels unevenly, and relationships emerge that no formal program architecture anticipated.

We are dealing, then, with a system that is bounded without being fully predictable, and continually moving without being unconstrained.

Where boundaries are insufficiently permeable, governance may become rigid. Information does not travel, relationships end at administrative boundaries, professional judgement is constrained, feedback is blocked, and adaptations are displaced elsewhere. This does not stop systemic movement; it can make its consequences less visible and less governable.

Where boundaries fail to preserve necessary distinctions, a different governance difficulty arises. Roles and accountability blur, information becomes noise, responsibilities can no longer be reliably located, and safeguards may be weakened. These are possible consequences of institutional arrangements, not opposite quantities on a measurable chaos scale.

Reform language can make this tension harder to recognise. Flexibility, integration, collaboration and information sharing are often presented as desirable in themselves, although each changes how boundaries operate. Their value depends on what a particular boundary exists to protect and what needs to move across it.

Permeability must itself be governed.

Permeability concerns what can cross a boundary while the boundary continues to perform its governing function. It is necessarily selective: relevant information may need to move while privacy is protected, relationships may need to continue while financial accountability remains bounded, and professional discretion may need room to operate within firm safeguarding requirements.

The practical judgement is therefore about what a boundary should permit, what it should protect, and the purposes and conditions governing each choice.

Figure 2 illustrates the risks at either extreme. Excessive rigidity can obstruct necessary adaptation, while the loss of essential boundaries can weaken coherence and accountability.




Figure 2. *Responding to systemic movement: preserving boundaries while allowing adaptation*

This also returns us to dynamic equilibrium, because the relationships among boundaries, movement and adaptation are themselves changing. Adaptation may strengthen the pursuit of a system's purpose, but it may equally help to sustain an unsatisfactory pattern of outcomes.

The governance question is this: is the system adapting in ways that improve its capacity to achieve its purpose, or only adapting enough to preserve its existing equilibrium?

The distinction can be brought together in the following proposition.

Bounded chaos names the constrained but continually changing movement of a complex institutional system. Governance does not create or maintain that movement; it can shape the boundaries, responsibilities and channels through which its consequences are recognised and addressed.

Stacey's examination of the tension between excessive control and the loss of necessary order provides an intellectual starting point for this proposition.^{10,14} Here, bounded chaos names the movement of actors and relationships within and across bounded institutions; the control–anarchy tension helps explain the governance challenge posed by that movement. These are not two separate definitions. Governance must preserve enough structure for purpose and accountability while allowing discretion, information exchange and adaptation. Feedback can reveal when controls obstruct necessary exchange or when insufficient structure weakens accountability, but only actors with the capacity and authority to respond can turn that information into action.

The term bounded is as important here as chaos. Rather than seeking to remove boundaries or regulate the quantity of movement, governance needs to consider what boundaries protect, what they permit to

cross and what consequences extend beyond them. The Fishbowl Effect makes those questions more tangible.

6. The Fishbowl Effect

Complex systems become administratively governable in part by being divided into bounded spaces. Departments have portfolios, programs have eligibility rules, organisations have contractual responsibilities, and funding streams have defined purposes. Performance frameworks and data systems likewise establish the populations, periods and activities that can be observed.

These divisions are necessary because authority, resources and accountability must be located somewhere. The difficulty arises when administrative boundaries fail to correspond to the ways outcomes are actually produced and experienced.

A person's employment trajectory can cross education, income support, employment services, training, work experience, recruitment and employment itself. It can also be shaped by transport, housing, health, caring responsibilities, family circumstances and technology. The trajectory continues while institutional responsibility repeatedly begins and ends around it.

We refer to the resulting separation between bounded institutional performance and longer trajectories as the Fishbowl Effect.

The Fishbowl Effect occurs when actors optimise within bounded parts of a system while the trajectories they are trying to influence extend beyond those boundaries.

The Fishbowl Effect concerns the limits of institutional visibility and accountability within a system whose interactions continue across those limits. Bounded chaos is not contained by the administrative divisions through which organisations make their work manageable; consequences continue to travel, whether or not the formal architecture recognises them.

The metaphor is apt because a fishbowl creates both visibility and constraint. From inside a bounded program or organisation, some aspects of performance are highly visible. Beyond the glass, visibility fades. A provider may know that a placement occurred but not what happened after its formal responsibility ended. A training organisation may know that a qualification was completed but not whether it led to employment. An employer sees the recruitment interaction but not the service journey that preceded it. A department sees outcomes within its own portfolio while the consequences surface elsewhere.

None of these actors needs to be irrational, unethical or incompetent for the Fishbowl Effect to occur. The problem is as much structural as behavioural.^{13,15}

Institutional boundaries constrain authority, but they do not necessarily constrain consequences. An adaptation that appears successful within one program may shift costs, obligations or risks to another institution or to the person whose trajectory extends across both. The actor who experiences or observes a problem may lack the authority to alter the conditions producing it.

Performance inside a fishbowl can therefore diverge from outcomes across the wider trajectory. A program may satisfy every requirement of its performance framework even as the person repeatedly crosses boundaries where knowledge, relationships or support are lost.

Measurement can reinforce this separation because information is often organised around the same funding and accountability boundaries as programs. Administrative records may show a sequence of successful events without revealing whether those events contribute to a coherent and sustainable trajectory.

The journey does not reset when responsibility changes. Knowledge gathered in one part of the system may remain valuable in another. Relationships can retain their value long after the activity that funded them has ended. Effective adjustments, capabilities and preferences do not become irrelevant because a program has closed.

Where boundaries are insufficiently permeable, the consequences frequently appear as repetition. People retell their histories, assessments are repeated, relationships have to be rebuilt, and employers are approached afresh. Knowledge gained earlier may need to be rediscovered simply because it could not travel across the boundary.

Removing the boundaries altogether would not solve this problem, since privacy, professional responsibility, financial accountability, safeguarding and authority still require protection.

A more useful response is selective permeability, through which relevant information can travel, relationships can persist where they matter, and support can taper or transfer without losing valuable learning. Responsibility may sometimes need to be handed on rather than merely brought to an administrative end.

Integration does not necessarily solve the Fishbowl Effect. Combining organisations or programs may simply create a larger fishbowl, and a larger administrative boundary still has an inside and an outside.¹⁵

The fishbowls are not independent containers. Their occupants make decisions, form relationships and respond to one another, and these interactions collectively constitute the wider system. It is therefore misleading to treat “the system” as wholly external to the people and institutions operating within it. Their individual authority may be limited, but their capacity to act together can extend beyond what any one actor can change alone. This does not erase unequal power or formal decision rights; it makes the relationships between them a matter for governance.

The central question is less about the number or size of the fishbowls than about whether their boundaries allow the exchanges on which outcomes beyond any one institution depend.

This becomes especially important where long-term outcomes are distributed across time and institutions. Decisions that are rational locally can combine to produce poor consequences at the level of the system. Where the necessary exchange across boundaries depends on goodwill, on individual practitioners or on informal workarounds, it is likely to be inconsistent and hard to observe.

The Fishbowl Effect therefore raises three practical questions. What must remain bounded? What must be able to cross the boundary? And what happens when something necessary to the wider trajectory can do neither?

Where an activity or responsibility falls outside the formal architecture, its underlying function may nevertheless remain necessary. The work can then move elsewhere, even though the architecture records the original activity as complete.

7. When necessary functions escape the architecture

Formal systems are generally designed around identifiable activities, responsibilities and transactions. A provider is funded to deliver particular services. A practitioner performs specified activities. A participant becomes eligible for assistance under defined conditions. An employer interacts with the system through particular programs or mechanisms. Payments are attached to recognised activities, milestones or outcomes.

Such specification is necessary for administration: governments need to know what they are commissioning, while organisations need a basis for allocating resources and assigning responsibilities. Activities provide tangible units through which funding, obligations and performance can be described.

It is important, however, not to confuse an activity with the function it performs.

An activity describes something that is done. A function describes what that activity accomplishes within the wider process through which an outcome is produced. Different activities can perform the same function, and a single activity can perform several functions at once.

Our work with UN Tourism on education-to-employment pathways offers a useful illustration. Across different national and institutional arrangements, the ways in which education providers engage employers, support transitions and learn from graduate outcomes vary considerably. The underlying functions, however, remain recognisable. This is a conceptual comparison, not evidence that any particular configuration is effective; identifying a function says nothing about how well it is performed. What it does allow is comparison across systems without assuming that they must all adopt the same institutional architecture.¹⁶

The distinction becomes important when systems are redesigned.⁴

A reform may remove, consolidate or replace an activity because it appears duplicative, inefficient or unnecessary. The activity can disappear from a program specification, a funding model or a service sequence. If the function it was performing is still needed, though, the need does not disappear with the administrative category through which it used to be recognised.

The disappearance of an administrative category does not, by itself, remove the practical need it was addressing.

Consider continuity. A system may not formally commission anyone to maintain continuity across successive services, yet continuity may still be essential if knowledge, relationships or support are to survive transitions. Coordination may not appear as a distinct funded activity, yet someone may need to coordinate the interactions among employers, services and participants. Translation between administrative requirements and individual circumstances may have no recognised program category, yet practitioners do it routinely, because services cannot otherwise be applied sensibly in practice.

These functions can be difficult to recognise precisely because they are embedded in activities described for other administrative purposes.

A practitioner described as providing case management may in fact be maintaining continuity across several systems. An employment consultant undertaking employer engagement may be negotiating job design, translating capability into employer requirements and sustaining a relationship that extends well beyond a single vacancy. A family member may be carrying information between services, coordinating appointments and explaining changes in circumstances. A participant may repeatedly reconstruct their own history because no formal mechanism allows relevant knowledge to travel with them.²

The formal architecture may record the activity, while the wider system continues to depend on the function beneath it.

When a necessary function is not adequately recognised, funded or governed within the formal architecture, it can be pushed beyond the boundary within which the system expects the outcome to be produced. We describe this as functional displacement.¹⁷

Functional displacement occurs when work necessary to the functioning of a system is still required but is pushed beyond the boundary within which it is formally recognised, resourced or governed.

Displacement takes several forms. A function can move from one organisation to another, from formal service delivery to informal practice, from paid to unpaid work, or from an institution to the participant or family. It can also become distributed across several actors, none of whom carries full responsibility for making sure it happens.

In these circumstances, the function may persist even though responsibility for performing it has shifted.

This matters because conventional performance information may not reveal the movement. Administrative systems tend to record recognised activities and transactions. Once a function is performed outside those categories, the system may have little capacity to see how much work is involved, who is doing it, what it costs or whether it is being done consistently.⁶

The resulting difference between what is recorded and what is actually required has important consequences for how reform is assessed.

The formal system can record the saving from removing an activity without recording the cost of relocating its function.

That cost may surface elsewhere in government expenditure. It may be absorbed by another organisation, added to practitioners' workloads or carried by employers. It may be transferred to participants and families. Sometimes it does not appear as expenditure at all, but as delay, repetition, frustration, reduced participation or an outcome that does not last.

A reform may consequently appear efficient from within the original fishbowl, while its costs look rather different when viewed across the wider trajectory.

Functional displacement also helps explain a familiar feature of service systems: the persistence of activities that seem inconsistent with formal program design.

Practitioners often develop ways of working that the formal architecture does not clearly recognise. They maintain contact after an administrative exit, spend extra time coordinating with another service, reinterpret an allowable activity to meet an emerging need, preserve an employer relationship across funding periods, or take on work that sits ambiguously between organisational responsibilities.

We use shadow delivery to describe this practitioner-level expression of functional displacement.

Shadow delivery does not imply concealment or misconduct. The term describes necessary work that occurs in the shadow of the formal delivery model. It is work that practitioners or organisations undertake because they judge it necessary to achieve or sustain the intended outcome, but which is poorly represented in the architecture through which delivery is commissioned and observed.

The qualification is important. Some informal activity does represent inappropriate practice, inefficiency or failure to comply with program requirements. Functional displacement should not become a blanket explanation for everything that happens outside formal specifications.

The relevant question is whether the activity performs a function on which the intended outcome depends. If it does, its persistence may be telling us something the architecture has failed to recognise.

Shadow delivery can therefore be a valuable source of diagnostic information. Repeated workarounds across practitioners or organisations may indicate that a necessary function has no adequate formal home. What looks like individual practitioner discretion may, taken together, reveal a structural gap.

This distinction also complicates how implementation fidelity is interpreted.

Variation from a prescribed model can reflect poor implementation. It can also arise because the formal model does not contain all the functions needed to respond to the environment in which it operates. In the first case, tightening compliance may improve performance. In the second, it may remove the very adaptations through which the system is currently managing to function.

The presence of variation alone cannot settle which explanation is more appropriate; it is necessary to understand the function that the variation performs.

Here functional displacement connects directly with the stable underperformance described earlier. Informal compensatory work can prevent a poorly configured system from failing outright. Practitioners bridge gaps, families maintain continuity, participants carry information and organisations absorb activities that fall between contracts. These adaptations can create enough permeability for the system to remain viable.

They can also make its structural weaknesses harder to see. A system can become dependent on work it does not formally recognise.

This presents reform with a paradox. Removing informal workarounds without first understanding the functions they perform can make performance worse. Yet leaving essential functions dependent on informal activity produces inconsistency, inequity and fragility.

Rather than seeking immediately to eliminate shadow delivery, governance first needs to make the underlying function visible and understand why it persists.

Once the function is visible, a governance decision becomes possible. The function might need to be formally commissioned. It might belong within an existing role. It might require a mechanism for transfer across organisational boundaries, or it might be better performed elsewhere. It may turn out to be unnecessary once another part of the architecture is changed. In some circumstances, flexible practitioner discretion may remain the best way of performing it.

The essential point is that the decision should concern the function, not merely the administrative label attached to whichever activity happens to be performing it.

The implications extend beyond the individual service activity to the way system redesign is conceived.

Reform commonly changes the allocation of activities. Programs are combined and services retendered. Payment structures change. Responsibilities move between organisations or levels of government. Particular forms of assistance are removed or replaced.

A functional perspective invites a further question about the consequences of each change.

Which functions must continue to occur if the intended outcome is to remain possible, and where will responsibility for those functions sit after the redesign?

That question can expose consequences an activity-based map would miss. It can also show why apparently separate parts of a person's journey are connected. Continuity, relationship maintenance, information transfer, capability development and adaptation often operate across several administrative events rather than belonging to any one of them.

A map limited to formal activities tends to lose these connections at administrative boundaries, whereas a functional perspective makes them easier to follow.

This brings us back to the problem of time. A function may begin in one program, matter in another and reveal its significance only much later. An activity can be completed and recorded while the function it was meant to perform remains unfinished. Conversely, a relatively small activity at one point can change what becomes possible later.

Understanding performance thus involves examining not only whether an activity occurred but also how it contributed to developments over time.

The discussion moves from the administrative event towards the longer trajectory in which its significance becomes apparent.

8. From events to trajectories

Administrative systems are understandably organised around identifiable events, because these are the points at which decisions and transactions can be recorded.

A person enters a service, completes an assessment, starts training, takes part in work experience, obtains employment, reaches a retention milestone or exits assistance. Each event can be defined, dated, recorded and, where appropriate, attached to a payment or a performance measure.

Events make complex activity governable. They provide identifiable points at which eligibility can be determined, resources allocated, responsibilities established and progress assessed. They also allow governments and organisations to aggregate individual activity into information about how programs are performing.

There is nothing inherently problematic in this. Many important outcomes are events, and even where an outcome extends over time, events provide necessary markers of progress. The difficulty arises when the event becomes a substitute for the outcome it is meant to represent.

Employment is a clear example. Obtaining a job is an important event. For many people it is the immediate objective of employment assistance, and an entirely legitimate point at which to recognise progress. But employment is also a relationship that unfolds over time.

After commencement, the worker has to enter the workplace, learn the role, build relationships, meet expectations and respond to changing circumstances. The employer adapts as well. Tasks change. Hours increase or decrease. Supervisors move on. Technology alters the work, and business conditions shift. Support needs emerge or recede, new opportunities appear and personal circumstances change.

The significance of the original employment event becomes clearer only as the relationship develops and its consequences unfold.

A job commencement may become the start of sustained employment and career development. It may become a short period of work followed by unemployment. It may lead to another job with a different employer. It may reveal capabilities or preferences that reshape later choices. It may end quickly and still contribute something important to a later employment outcome.

An event cannot tell us all of this at the moment it is recorded, which is why it is useful to distinguish employment events from employment trajectories.

A trajectory describes the direction and pattern produced by a sequence of events and changing conditions over time. It can include entry to employment, retention, interruption, progression, changes of job, periods outside work, changes in hours or occupation, the development of capability and reconnection to employment after disruption. A trajectory need not be linear to be successful.

The qualification matters. Administrative models often represent progress, implicitly, as a sequence of increasingly desirable states: preparation leads to placement, placement to retention, and retention to independence from assistance. For some people, employment does roughly follow this sequence. For others, sustainable participation includes interruption, experimentation, changes of employer, periods of additional support, retraining or movement between different forms of work.

A change in direction need not indicate that a trajectory has failed, just as an apparently unchanged administrative status does not necessarily demonstrate success.

A person may remain employed while their hours decline, the quality of their job deteriorates, their opportunities narrow or they come to depend increasingly on informal support. A stable administrative

state can conceal a changing trajectory, just as a changing sequence of events can conceal longer-term progress.

Thinking in terms of trajectories also alters the way continuity is understood.

Continuity does not require one provider, program or practitioner to remain involved indefinitely. Nor does it require the person to stay in the same job. Continuity concerns what needs to persist as the structures around the person change.

Relevant knowledge may need to persist. Relationships may retain their value. An understanding of capability, preferences or effective adjustments may remain useful. Connections with employers or industries may matter later. Support may need to be able to return when circumstances change, and learning from earlier experience may need to inform later decisions.

Although the institutions and people involved may change, some of the functions supporting the trajectory may need to continue.

This is where the Fishbowl Effect becomes especially relevant.

Programs have beginnings and ends because programs need boundaries. Organisations accept and discharge responsibility because responsibility cannot be unlimited. Funding periods close, contracts expire, eligibility changes and employers come and go.

A trajectory continues across the boundaries at which individual programs and responsibilities begin and end. What appears to one organisation as a completed exit may, for the person concerned, be another transition in a much longer process.

This creates a potential mismatch between the events systems are commissioned to deliver and the longer outcomes they ultimately seek.

Governments can commission events even when the outcome they ultimately seek is a trajectory.

A placement payment purchases or rewards an observable event. A retention milestone extends observation a little further in time. Training completion records another event. Each may be useful and appropriate.

No single event, however, captures whether a person's relationship with employment is becoming more sustainable, more productive, more resilient or better able to support future progression.

Adding more milestones does not necessarily solve the problem. A sequence of snapshots remains a sequence of snapshots unless the relationships between them are examined.^{4,5}

This matters because administrative performance can look strong while trajectories remain unstable. A person may generate several successful employment commencements over several years while repeatedly returning to unemployment. Each commencement may satisfy the relevant program measure. Viewed over time, though, the pattern raises a different question: what is preventing employment from becoming durable?

The reverse is also possible. A period with few reportable outcomes may belong to a trajectory in which capability, confidence, employer relationships or vocational direction are developing in ways that later prove significant.

A trajectory perspective changes the question from “Did the event occur?” to “What did the event change?”

Did training change what employment became possible? Did a placement lead to stronger attachment to work? Did a period of employment build capability that carried over into a later role? Did support reduce because the employment relationship had become more sustainable, or because formal responsibility had ended? Did a job loss return the person to the beginning of the system, or did the knowledge and capability gained in previous employment persist?

The organisation responsible for the original event cannot always answer these questions. That is precisely the point. If outcomes extend beyond the boundaries of individual programs and organisations, a system concerned with those outcomes needs some capacity to observe beyond those boundaries.

This does not mean every provider should be held accountable for every subsequent event in a person's life. Such an approach would be neither fair nor administratively workable. Accountability should reflect what actors can reasonably influence. Observation and accountability, however, are not the same thing.

A system can seek to understand what happens after a program ends without making the former provider responsible for everything that follows. Without such observation, in fact, the system has little capacity to learn whether earlier interventions contributed to durable change or merely produced temporary administrative success.

The distinction is especially important in employment systems, where responsibility is often distributed. An education provider may contribute capability. An employment service may assist with entry to work. An employer provides the environment in which employment takes place, and other services may influence whether that employment remains viable. The person is the only actor necessarily present across the whole trajectory.

Employment outcomes are produced through relationships extending over time, rather than being contained within any single institutional boundary.

A trajectory perspective also brings transitions into focus. When a system is organised mainly around bounded interventions, the points between those interventions can receive less attention than the activities within them. Yet transitions are often exactly where information, relationships and responsibility are most vulnerable to loss.

A person moving from education into employment assistance, from employment assistance into work, from short-term support into ongoing employment, or from one employer to another may cross several administrative boundaries in a short period.

At each transition the system faces the permeability question developed earlier. What must move with the person? What can appropriately end? What must be able to return if circumstances change?

These transitions bring us to a particularly important question about what continuity requires.

What needs to persist when the organisations, programs, employers, funding arrangements and circumstances around the person change?

The answer will differ between people and contexts. It may include knowledge, relationships, capability, adjustments, access to support, vocational direction, or mechanisms through which assistance can be reactivated. What matters is that continuity is not confused with permanence.

A trajectory perspective does not require services to retain people indefinitely; rather, it asks that things of continuing value should not be discarded whenever an administrative boundary is crossed.

This gives the idea of trajectory a relevance well beyond employment services. Education-to-work transitions, career progression, workforce re-entry and responses to displacement all unfold across time and institutional boundaries. Similar questions arise wherever the outcome being pursued develops over a longer period than the intervention commissioned to influence it.

The event remains useful, but its significance becomes clearer when examined in relation to what came before it, what followed and what changed as a result.

None of this removes the need for snapshots. Nor does it imply that every individual journey must be predicted or centrally managed. The argument of this paper is precisely that future conditions cannot be fully known. Trajectory thinking instead creates a different requirement for observation.

For this kind of understanding to inform governance, information about consequences must be able to find its way back. This leads to the role of feedback as system infrastructure.

9. Feedback as system infrastructure

Because the consequences of intervention cannot be fully specified in advance, systems need ways to learn from what happens after decisions have been made and implemented.

Feedback therefore matters not only as a quality-improvement technique but also as part of the infrastructure that enables a dynamic system to respond to changing conditions.

Employment and workforce systems already collect a great deal of information. Participants are assessed, activities recorded, payments processed and outcomes reported. Employers provide information. Providers monitor performance. Governments analyse administrative and labour-market data, and evaluations examine the effects of programs and reforms.

Yet the presence of extensive information does not necessarily mean that consequential feedback is occurring.

A system can be data-rich and learning-poor.

Information becomes feedback when it can influence the conditions that helped produce what has been observed. That takes more than collection. The information must be relevant enough to reveal a consequence, timely enough to remain useful, able to travel across the necessary boundaries, and able to reach an actor with the authority or capacity to respond. ^{1,18}

Systems become adaptive when information about consequences can travel back to a point where the conditions producing those consequences can be changed.

Feedback communicates consequences but does not itself bring about change. Information may be overlooked, misinterpreted or received by someone unable to act on it, so an effective response depends on reaching a person, institution or coordinated group with relevant authority, capacity and willingness. Even then, the response remains provisional until its effects can be examined.

The destination of feedback matters as much as the information being transmitted.

Feedback does not necessarily need to return to the actor responsible for the original intervention. Where outcomes are distributed across organisations and through time, the actor who first observes a consequence may neither have caused it nor have the authority to address it.

The relationship between bounded chaos and the Fishbowl Effect makes the location of corrective authority particularly important. Consequences may be experienced or observed outside the institution whose decisions helped produce them. The greater the separation between the site of consequence and the locus of authority, the more demanding the requirements for feedback that can support a corrective response. This is a proposition for investigation, not an established causal law.

This gives rise to an important distinction between observation, accountability and intervention. ¹⁹

An actor may observe something without being accountable for the circumstances that produced it. An organisation may recognise a problem without having the responsibility or authority to intervene. Conversely, an actor with the authority to change a condition may have little direct sight of its consequences.

These distinctions arise legitimately in systems where authority and accountability are bounded; they do not imply that everyone who observes a consequence should become responsible for correcting it.

That approach would weaken the boundaries through which responsibilities are allocated and would invite exactly the dissolution discussed earlier. Practitioners, providers, employers and government agencies need limits to their authority as well as clarity about their obligations.

Between observing a consequence and having responsibility to intervene lies the important work of communicating what has been learned.

An actor may reasonably conclude: “I can see what has happened, but it is not my responsibility to intervene.” That conclusion does not mean the information should stop there.

If what has been observed is consequential for the person’s trajectory, or reveals something important about how the system is functioning, the governance question becomes whether there is a pathway through which that information can reach someone able and authorised to respond. Figure 3 sets out the sequence this implies.¹⁹



Figure 3. *From observation to authorised response*

Not every observation should set this sequence in motion. Systems already produce more information than anyone can meaningfully process, and effective feedback depends on discrimination. The aim is not a system in which every piece of information travels everywhere.

That would be information permeability without governance. What the system needs instead are mechanisms through which consequential information can cross appropriate boundaries.^{15,19}

What counts as consequential depends on the purpose and context of the system. It might be evidence that apparently successful employment outcomes repeatedly fail once formal support ends. It might be information that a particular transition between services consistently produces a loss of continuity. It might show that employers meet the same problem across multiple providers, that participants are repeatedly and unnecessarily reassessed, or that practitioners in different organisations are all performing the same unfunded coordination function. Any single observation may seem minor. Patterns become visible when observations can accumulate.

This is another reason feedback cannot be reduced to the responsibility of individual practitioners. A practitioner may recognise a problem but reasonably lack the authority to change the funding, program or organisational condition producing it. A provider may identify a recurring issue but be unable to alter a contractual requirement. A commissioner may have the authority to change that requirement but little sight of what happens in everyday practice.

Where information and authority reside in different parts of the system, feedback infrastructure provides the connections through which one can inform the other.

Feedback moving from practice to governance

An applied illustration arises from CDERP's evaluation of the STRIDE program delivered with Devonfield Enterprises in north-west Tasmania. The evaluation considered participant experience, practitioner capability, practice fidelity and organisational quality-assurance evidence. It identified substantial development in practitioner competence and implementation, while also identifying areas requiring further attention at management and governance levels. In particular, the report recommended stronger leadership capability in fidelity oversight, better integration of participant and practice evidence into strategic reporting, and more systematic governance of employer relationships.

The significance is not simply that the evaluation collected several kinds of evidence. Information arising from direct practice exposed organisational conditions that individual practitioners could not address alone. The findings could therefore inform decisions about supervision, reporting, capability and the continuity of employer engagement. In this sense, the case illustrates the potential movement of feedback from practice to governance. It does not establish that the recommended actions were subsequently implemented or that they improved outcomes; verifying those responses would be necessary to demonstrate a completed corrective feedback cycle.²⁰

This connection matters most in systems shaped by the Fishbowl Effect. Information about what happens beyond an administrative boundary may never return to the actor whose earlier decisions influenced the trajectory. The provider knows that employment began but not why it ended. The training organisation knows that a qualification was completed but not whether it led to work. The commissioner knows that a milestone was reached but not whether the underlying outcome lasted.

The difficulty in such cases may lie less in a shortage of data than in the interruption of the pathway through which consequences could inform later decisions.

Observation, accountability and intervention perform different roles: the first makes consequences visible, the second allocates responsibility, and the third involves the authority and obligation to act. Feedback can connect these roles without requiring them to be collapsed into one.

This distinction becomes especially significant when similar consequences are observed repeatedly across different parts of the system.

If one practitioner sees a problem that lies outside their responsibility, it may be enough to refer it, communicate it or simply record it. If many practitioners, organisations or participants keep encountering the same problem, and no actor has clear responsibility for the underlying function, the issue changes character.

At that point, what appeared to be a series of individual operational difficulties may warrant attention as a governance problem.

If everyone can see a recurring problem but nobody owns the function required to respond to it, the absence of ownership is itself information about the system.

This is also where feedback becomes connected to functional displacement.

A missing feedback pathway allows displaced functions to stay invisible. Practitioners compensate locally. Families bridge gaps. Participants carry information. Providers improvise. Because the immediate problem keeps being resolved within individual fishbowls, the wider system learns little about the structural condition that generates the need for compensation.

Successful local adaptation can weaken learning at the level of the system. The compensation that keeps an underperforming system resilient also deprives governance of the signal that something needs to change. Feedback must therefore travel not only horizontally, across operational boundaries, but vertically, towards governance.

A recurring workaround may need to become visible to organisational management. A recurring organisational problem may need to become visible to a commissioner. A pattern spanning several commissioned systems may need to become visible at the level of policy. Consequences that cross employment, education, income support or other domains may need to be observed at a level above any single program.

This does not mean creating a single central authority responsible for the entire system. It means that feedback pathways should reflect the level at which the relevant condition can actually be changed. A local practice problem should not require system reform. A structural commissioning problem cannot be solved by practitioner effort alone. The right destination for feedback depends on where the mechanism producing the consequence is located.

The same reasoning applies when actors describe a failing program as the product of “the system” and conclude that they cannot change it. Some constraints genuinely require decisions beyond their authority. Others persist through the practices, assumptions and relationships of the actors themselves. The practical question is therefore not only who has formal authority, but which actors can work in concert to bring evidence to that authority, exercise existing discretion or change the relationships that reproduce the problem. Collective agency does not guarantee reform, but it challenges the assumption that bounded individual authority amounts to collective powerlessness.


There is a further dimension to feedback in a system whose relationships and conditions are continually changing.

Much quality-assurance thinking imagines a loop in which information travels backwards: an intervention occurs, its consequences are observed, the learning returns to the original actor and practice is adjusted. Dynamic systems require something more complex.


Information may need to travel backwards, sideways and upwards. It may need to cross organisational boundaries, move between operational and strategic levels, or reach a different part of the system altogether. The person best placed to observe a consequence may be far removed from the actor able to change its cause.

Feedback may be better understood as a network of learning pathways than as a single loop returning information to its original source.

Even then, the familiar image of the feedback loop has a further limitation. A conventional loop implies a return to the starting point (Figure 4, left below).



Closed loop



Moving spiral

Figure 4. *The closed loop and the moving spiral*

The image of a return to the original starting point is nevertheless misleading in a system that has already changed.

By the time information is received and interpreted, the original action has altered conditions, actors have responded, relationships may have shifted and new opportunities or constraints may have emerged. Adaptation therefore takes place in a setting that differs from the one in which the intervention began.

In this sense, the feedback loop itself moves with the system.

It is better represented as a spiral (Figure 4, right), or as repeated movement through a changing maze, in which each cycle of action and learning begins from a somewhat different position.

This continuing movement is central to the importance of feedback infrastructure for dynamic equilibrium.

Its purpose is not to return the system to some previous state, nor to eliminate every source of uncertainty. Its purpose is to give the system enough capacity to detect consequential change and respond before changing conditions make its existing assumptions ineffective.

Data systems are only one part of that infrastructure, alongside selective permeability, interpretation, authority and opportunities for learning.

A technically sophisticated information system that produces reports no one uses is not effective feedback infrastructure. Nor is a participant survey whose findings cannot influence commissioning, a practitioner escalation process that keeps ending at an organisational boundary, or an evaluation completed after the relevant policy settings have already changed.

The more useful test is whether information about consequences can reach a place where the conditions producing them might be changed, rather than simply whether it was collected.

The argument can now be brought together. A blueprint establishes architecture, but intervention also changes behaviour; events have consequences that unfold across trajectories, and feedback makes some of those consequences available to the wider system. How governance responds to what it learns is therefore as important as the design with which reform began.

This leaves the question of how systems can respond intelligently without assuming that everything occurring within them can be controlled. It is the practical challenge of governing the moving maze.

10. Governing the moving maze

The argument developed so far should not be mistaken for a case against system design.

Governments still have to make choices. They must establish purposes, determine eligibility, allocate public resources, protect rights, distribute authority, create accountability and decide which outcomes matter. Organisations need structures through which responsibilities can be exercised. Participants and employers need enough certainty to understand what they can expect from the systems they deal with.

Architecture remains necessary, although it cannot by itself ensure that a system continues to pursue its purpose under changing conditions.

A system designed at one point in time will meet circumstances that were not fully known when it was designed. Actors will respond to its incentives and requirements. Technologies will change, economic conditions will shift and organisations will adapt. Informal practices will emerge. Some anticipated relationships will not develop as expected, while others will become important in ways that could not have been specified in advance.

The task of governance therefore extends beyond deciding what the system should look like. It includes maintaining the conditions through which the system can keep pursuing its purpose as the environment around and within it changes.

This is a different conception of governance.^{13,18}

It does not ask government to predict every future circumstance. Nor does it call for continuous redesign in response to every variation in performance. Constant intervention can itself create instability, particularly when actors have too little time to adapt to one set of conditions before the next arrives. Governing a dynamic system instead requires judgement about what should remain relatively stable and what should be able to change.

Some elements of a system need persistence. Purpose should not change simply because circumstances become difficult. Rights should not depend on administrative convenience. Accountability should remain identifiable. Essential functions should continue even when the organisations or activities through which they are delivered change.

Other elements need to be more adaptable. Operational practice may need to respond to individual circumstances. Information pathways may need to change as new consequences come into view. Relationships may need to continue across administrative transitions. Resources may need to be redirected as demand changes, and commissioning arrangements may need to evolve when evidence shows that necessary functions have been displaced.

The challenge is not to choose between stability and adaptation. It is to decide what should be stable, what should be adaptable, and how the relationship between the two should be governed. The following seven conditions make that governance responsibility explicit, and a closing synthesis draws them together. They are interdependent requirements for continuing judgement, not a sequential checklist or a blueprint for an ideal system.

Condition 1. Preserve clarity of purpose

The first condition concerns the continuing clarity of the system's purpose.

Adaptation without purpose becomes drift. If organisations and practitioners are given flexibility without a clear sense of what the system exists to achieve, local adaptation can gradually carry activity away from the outcome originally intended.

Purpose provides an important boundary within which local judgement and adaptation can occur.

This is especially important where performance measures have become proxies for broader objectives. A placement, a training completion or a retention milestone may represent progress towards an employment objective, but the measure should not displace the purpose for which it was introduced.

A system can become steadily better at producing its measures while becoming worse at producing the outcome those measures were meant to represent.

Governing purpose means keeping the two distinct. It also means asking, from time to time, whether the activities and measures within the system are still connected to the purpose they are meant to serve. In a moving system, yesterday's reasonable proxy can become tomorrow's source of distortion.

Condition 2. Sustain essential functions

Purpose identifies what a system seeks to achieve, while functions describe the work needed to make that achievement possible.

The distinction allows governance to become less dependent on particular organisational forms. Programs can change. Providers can change. Funding mechanisms can change. Technologies can alter how services are delivered, and activities can be redesigned or replaced. Essential functions may nevertheless persist.

They should not, however, be treated as a permanently fixed inventory. Their significance depends on the purpose of the system and the conditions under which that purpose is pursued. Functional governance therefore requires periodic examination not only of how functions are performed, but of whether functions identified earlier are still necessary and whether new ones have emerged.

If continuity is necessary, it has to occur somewhere. If employers need effective engagement, some part of the system has to provide it. If information needs to travel across transitions, there must be a mechanism through which it can move. If participants need support to adapt to changing employment circumstances, responsibility for that support has to sit somewhere within the wider system.

This does not mean every function should be centrally prescribed. It means that redesign should not confuse removing an activity with removing the need for the function it performed.

Functional governance can begin with a question about what the system must continue to accomplish.

What work must continue to occur for the system's purpose to remain achievable, whatever the organisational form through which that work is currently delivered?

Uncertainty about the location and responsibility for essential functions creates conditions in which functional displacement is more likely.

Functional governance does not require uniform institutional architecture. It requires clarity about which functions are essential to the system's purpose, where responsibility for them sits, how their performance can be observed, and how arrangements can be adapted when changing conditions show that they are no longer being performed adequately. This permits comparison across different arrangements without presuming that a common structure, or the mere formal presence of a function, guarantees an effective outcome.

Condition 3. Govern boundaries and permeability

Boundaries remain essential to coherent governance.

The case for permeability should not become a case for removing them. Authority, accountability, privacy, safeguarding, eligibility, professional scope and the stewardship of resources all require boundaries of some kind.

The governance question concerns what those boundaries should allow to cross. A system should examine its boundaries not only as divisions of responsibility but as points of exchange.

Does relevant information travel? Can relationships persist where continuity matters? Can learning about consequences return? Can responsibility be transferred clearly rather than simply ended? Can people move between programs without repeatedly losing what was known about them before the transition?

And at the same time, are privacy, authority and accountability adequately protected?

These questions cannot be answered once for the whole system. Different boundaries perform different functions and call for different degrees of permeability.

The objective is not an integrated system in which every boundary disappears. It is a system whose boundaries are selectively permeable to the things the wider trajectory needs in order to function.

Condition 4. Observe trajectories without claiming ownership

A trajectory perspective raises a further challenge.

No government department, provider, employer or practitioner can reasonably own a person's entire employment trajectory. The person moves through circumstances and relationships that extend beyond the authority of any single institution.

Trying to assign complete responsibility for the trajectory to one actor would simply be a different form of overreach.

But the absence of complete ownership does not mean the trajectory should remain invisible. Governance can distinguish between owning a trajectory and being able to observe it.

A system can examine what happens after a program ends without holding the former provider responsible for every subsequent event. It can identify recurring failures at transitions without requiring one organisation to manage every transition. It can examine whether employment outcomes last without assuming that employment services alone determine how long they last.

Observing trajectories allows the system to learn across boundaries while keeping reasonable limits on accountability. This matters most where administrative success and longer-term outcomes diverge.

If a system keeps producing successful events but unstable trajectories, that pattern should be visible somewhere within its governance.

Condition 5. Connect feedback to actionable authority

Observation alone does not make a system adaptive. Information about consequences must be able to reach a point where something relevant can be changed. The distinction drawn earlier between observation, accountability and intervention is therefore central to governance.

An actor who sees a consequence does not automatically become responsible for correcting it. But a system in which consequential information repeatedly stops at the boundary of whoever observed it cannot readily learn.

Governance needs pathways through which consequential information can move to the appropriate level. Some information should stay local because the relevant response is operational. Other information should move to organisational management, to commissioning or to policy, because the mechanism producing the consequence sits at one of those levels.

This suggests a principle of feedback proximity: information should travel far enough through the system to reach a point that can influence the relevant condition, and need not travel further merely because it can.

The principle guards against two opposite failures. The first is the fishbowl, in which important information cannot escape its local boundary. The second is an indiscriminate information environment in which everything is escalated and the result is noise rather than learning.

Effective feedback depends on both permeability and discrimination.

Condition 6. Evaluate adaptation against purpose

Adaptation is necessary in a dynamic system, but it should not be romanticised.

Practitioners can adapt around poor processes. Families can compensate for missing functions. Organisations can learn to succeed within distorted incentives. Providers can become highly effective at producing administratively valued events without changing the trajectories those events are meant to influence.

The presence of adaptation tells us that the system is responsive. It does not tell us that the response is desirable. Governance needs some capacity to distinguish adaptation that advances purpose from adaptation that merely preserves viability. This is where stable underperformance becomes especially important.

A system may keep operating because its actors are remarkably good at compensating for weaknesses in its architecture. Removing those adaptations without understanding them can cause harm. Leaving them permanently informal can entrench inconsistency and conceal structural problems.

The governing task is first to make the adaptation visible, then to ask what function it performs and why the formal system needs it to occur.

Sometimes the right response will be to formalise the function. Sometimes it will be to change the condition that creates the need for it. Sometimes the adaptation should remain local and discretionary.

The answer cannot be read off from the bare fact that variation exists.

Condition 7. Examine the effects of intervention

Perhaps the hardest implication of the moving-maze perspective is that governance must treat intervention itself as capable of generating system effects.

Policy does not enter an inert environment. Actors anticipate it, interpret it and respond to it. New performance measures redirect attention. New funding models change organisational behaviour. New eligibility rules shift demand elsewhere. New accountability requirements change what information is collected and which activities receive priority.

The effects of a reform include not only the outcomes deliberately designed into it but also the responses its introduction provokes.

Implementation, then, should not be understood simply as the period during which the intended design is progressively put in place. It is also a period of observation.

What behaviours are emerging? Which functions are moving? What workarounds are appearing? Where are new boundaries blocking necessary exchange? Which assumptions built into the design are proving robust, and which are meeting conditions that nobody anticipated?

These observations should not automatically trigger further redesign. A dynamic system needs time to adapt, and excessive policy volatility can prevent stable relationships and practices from forming.

The governance challenge is to distinguish variation that represents normal adaptation from patterns that point to a structural problem requiring intervention.

That distinction calls for judgement, not a formula.

Closing synthesis: governing for dynamic equilibrium

Taken together, these conditions describe what it means to govern for dynamic equilibrium.

The objective is not to specify every future interaction. Nor is it to create a system so flexible that every actor can respond however they choose.

The task is to maintain the conditions within which the system can keep pursuing its purpose while learning from the consequences of its own operation.

That requires enough stability for rights, responsibilities, essential functions and accountability to remain meaningful. It requires enough permeability for information, relationships, people and learning to cross appropriate boundaries. It requires observation that extends beyond administrative events to trajectories. It requires feedback pathways through which consequential information can reach actors able to respond. And it requires enough adaptive capacity for the system to change when the assumptions on which its architecture was built no longer hold.

These seven conditions are not specifications for an ideal system. They are continuing governance obligations, each requiring attention as circumstances and relationships change.

Their institutional expression will differ across systems, jurisdictions and periods. Different purposes require different boundaries. Different risks justify different levels of control. Different environments call for different forms of adaptation.

The distinction that matters is between designing a system to remain unchanged and governing a system so that it can change without losing what must remain stable.

This brings us back to bounded chaos.

Governance maintains necessary boundaries. Interaction produces movement. Feedback communicates consequences. Actors with authority and capacity, sometimes acting in concert, can respond. The effects of their responses become part of the system's continuing movement. Dynamic equilibrium emerges from these changing relationships; it is not a state governance can fix in place.

The maze remains in motion. The task is not to stop it moving, nor to pretend that its future configuration can be perfectly mapped.

The task is to govern it while it moves.

11. An invitation rather than a blueprint

This paper has deliberately avoided offering a model of the ideal employment or workforce system. To do so would reproduce the very assumption it has set out to question.

Our argument has been that systems of this kind are dynamic, relational and bounded. They are continually altered by the behaviour of the actors within them, by changing external conditions, and by the interventions intended to improve them.

That does not make design irrelevant. It changes the task of design and of governance.

Architecture establishes boundaries, responsibilities and incentives. How a system functions, however, also depends on what can move across those boundaries, whether necessary functions remain visible, whether trajectories can be observed beyond administrative events, and whether information about consequences can reach the places where conditions can be changed.

The propositions set out here are offered as a contribution to that discussion, not as a completed theory of system governance. They draw together insights from CDERP's research, conceptual development, organisational analysis and applied engagement with employment and workforce systems.

They also open questions that need further examination.

What must remain bounded?

What must be able to move?

What must we be able to see?

And when the system changes in response to what we do, how will we know what to change next?

Statement on the Use of Artificial Intelligence

Artificial intelligence tools were used during the preparation of this paper to support editorial development, language refinement, structural review and the examination of conceptual clarity and consistency. The underlying research, conceptual propositions, interpretations and conclusions remain the responsibility of the author. AI assistance supported intellectual exploration and communication but did not substitute for scholarly judgement or authorship. The author retains responsibility for the accuracy, integrity and final content of the paper.

References

1. Meadows, D. H. (2008). *Thinking in systems: A primer* (D. Wright, Ed.). Chelsea Green Publishing.
2. Lipsky, M. (1980). *Street-level bureaucracy: Dilemmas of the individual in public services*. Russell Sage Foundation.
3. Sterman, J. D. (2000). *Business dynamics: Systems thinking and modeling for a complex world*. Irwin/McGraw-Hill.
4. Pawson, R., & Tilley, N. (1997). *Realistic evaluation*. SAGE Publications.
5. Campbell, D. T. (1979). Assessing the impact of planned social change. *Evaluation and Program Planning*, 2(1), 67–90. [https://doi.org/10.1016/0149-7189\(79\)90048-X](https://doi.org/10.1016/0149-7189(79)90048-X)
6. Bevan, G., & Hood, C. (2006). What's measured is what matters: Targets and gaming in the English public health care system. *Public Administration*, 84(3), 517–538. <https://doi.org/10.1111/j.1467-9299.2006.00600.x>
7. Hood, C. (2006). Gaming in Targetworld: The targets approach to managing British public services. *Public Administration Review*, 66(4), 515–521. <https://doi.org/10.1111/j.1540-6210.2006.00612.x>
8. Smith, P. (2018). *A systems analysis of factors that lead to the successful employment of people with a disability* [Doctoral thesis, University of Sydney]. Sydney eScholarship Repository. <https://ses.library.usyd.edu.au/handle/2123/18610>
9. Smith, P. (2026). *The Equilibrium Systems Model of Employment (ESME): A unique integration of people, providers, employers and policy for sustainable work*. Labour & Industry. <https://doi.org/10.1080/10301763.2026.2655000>
10. Stacey, R. D. (1995). The science of complexity: An alternative perspective for strategic change processes. *Strategic Management Journal*, 16(6), 477–495. <https://doi.org/10.1002/smj.4250160606>
11. Cilliers, P. (1998). *Complexity and postmodernism: Understanding complex systems*. Routledge.
12. Holling, C. S. (1973). Resilience and stability of ecological systems. *Annual Review of Ecology and Systematics*, 4, 1–23. <https://doi.org/10.1146/annurev.es.04.110173.000245>
13. Ostrom, E. (2010). Polycentric systems for coping with collective action and global environmental change. *Global Environmental Change*, 20(4), 550–557. <https://doi.org/10.1016/j.gloenvcha.2010.07.004>
14. Stacey, R. D. (1992). *Managing the unknowable: Strategic boundaries between order and chaos in organizations*. Jossey-Bass.
15. Ansell, C., & Gash, A. (2008). Collaborative governance in theory and practice. *Journal of Public Administration Research and Theory*, 18(4), 543–571. <https://doi.org/10.1093/jopart/mum032>

- 16.** Smith, P. (2026). Connecting tourism graduates with employment opportunities [Commissioned manuscript in preparation for UN Tourism; unpublished].
- 17.** Smith, P. (2026). Governing integrative employment functions in transactional employment systems [Manuscript under review, Australian Journal of Public Administration]. Centre for Disability Employment Research and Practice.
- 18.** Argyris, C., & Schön, D. A. (1978). Organizational learning: A theory of action perspective. Addison-Wesley.
- 19.** Bovens, M., Schillemans, T., & Goodin, R. E. (2014). Public accountability. In M. Bovens, R. E. Goodin, & T. Schillemans (Eds.), *The Oxford handbook of public accountability* (pp. 1–20). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780199641253.013.0012>
- 20.** Centre for Disability Employment Research and Practice (CDERP). (2026). Skilling Northwest Tasmania: STRIDE final report [Project evaluation report prepared in partnership with Devonfield Enterprises; delivered to the Department of Social Services, Canberra, May 2026].