# The Separation of Objective and Dynamics

Variational Quantum Algorithms · Module 1 — Hybrid Quantum Classical Systems

Node 1.9

## Position in Knowledge Graph

**System:** Variational Quantum Algorithms **Structural Domain:** Hybrid Quantum-Classical Systems **Node:** 1.9

This is the **closing node** of Module 1 and its most important structural claim.

The thesis of Module 1 compresses to a single sentence: **in a variational quantum algorithm, the objective and the dynamics are two different things, and most confusions in the field come from forgetting that.**

This node is the formal statement of that claim, and the justification for treating "objective" and "dynamics" as the two independent axes along which the rest of the course is organized.

---

## The Two Components

Every VQA has two separable pieces.

**The objective.** A function $C : \mathbb{R}^P \to \mathbb{R}$, defined implicitly by the choice of ansatz $U(\boldsymbol{\theta})$, Hamiltonian $H$, and hardware noise model. This function has a shape: local minima, saddle points, flat regions, periodicities. It has a noise profile: the variance of $\widehat{C}(\boldsymbol{\theta})$ as a function of $\boldsymbol{\theta}$ and shot budget. Once the ansatz and Hamiltonian are fixed, the objective is **completely determined** — it exists as a mathematical object independent of any optimizer.

**The dynamics.** A rule $\mathcal{A}$ for updating $\boldsymbol{\theta}_k \to \boldsymbol{\theta}_{k+1}$ given feedback from $\widehat{C}$. Examples: gradient descent, Adam, HOPSO, COBYLA, SPSA, Bayesian optimization, simulated annealing. The dynamics are **completely determined** by the choice of optimizer and its hyperparameters — they exist as an algorithm independent of any particular objective.

**The behavior of the VQA is the interaction of these two things.** Not the objective alone. Not the dynamics alone. Their composition.

---

## Why This Separation Is Not Obvious

It should be obvious. It is not. The reason is that in most of the VQA literature, discussions of "why VQA is hard" fold both components into a single narrative and then blame one of them.

Typical examples of the confusion:

- "VQE converges slowly because the landscape has barren plateaus." *Half true.* Barren plateaus are an objective property. But whether slowness becomes failure depends on the dynamics — an optimizer that averages over shots differently can cope; one that takes single-sample gradient steps cannot.

- "Adam doesn't work well for VQA." *Half true.* Adam has particular dynamical properties (momentum, adaptive learning rates per parameter) that interact in particular ways with particular objective shapes. On some objectives Adam is great; on others it is terrible. The statement "Adam doesn't work well for VQA" is the statement "averaged across the VQA objectives the community cares about, Adam's dynamics interact poorly," which is a much narrower claim than the plain reading suggests.
- "Parameterized circuit $X$ is expressive enough to solve problem $Y$." *Expressivity is an objective property.* But reaching the expressive floor requires dynamics that actually get you there. Expressivity bounds are upper bounds on what is reachable, not statements about what any particular optimizer will find.

The pattern in all of these: conflating a shape-property of the objective with an emergent-property of the optimizer-plus-objective system.

---

## The Formal Statement

Let $\mathcal{O}$ be an objective (cost function plus noise model) and let $\mathcal{A}$ be a dynamics (update rule plus hyperparameters). Define the **behavior** of the pair $(\mathcal{O}, \mathcal{A})$ as the distribution over trajectories $\{\boldsymbol{\theta}_k\}_{k=0}^{T}$ induced by running $\mathcal{A}$ on $\mathcal{O}$ from random initializations.

**Claim 1 (decomposition).** The behavior is not a function of $\mathcal{O}$ alone or $\mathcal{A}$ alone. It is a function of the pair.

**Claim 2 (non-trivial interaction).** There exist pairs $(\mathcal{O}_1, \mathcal{A}_1)$ and $(\mathcal{O}_2, \mathcal{A}_2)$ such that swapping dynamics — evaluating $(\mathcal{O}_1, \mathcal{A}_2)$ — produces qualitatively different behavior. In particular, convergence on $(\mathcal{O}_1, \mathcal{A}_1)$ does not imply convergence on $(\mathcal{O}_1, \mathcal{A}_2)$, and vice versa.

**Claim 3 (separable study).** Despite the non-trivial interaction, one can profitably study the two components separately. Properties of $\mathcal{O}$ (barren plateaus, landscape geometry, concentration phenomena) are well-defined without reference to any optimizer. Properties of $\mathcal{A}$ (stability, convergence rates, noise tolerance) are well-defined without reference to any particular objective. The composition happens last.

Claims 1 and 2 are empirical observations backed by literature (e.g., Arrasmith et al. on optimizer choice in VQA). Claim 3 is the organizational principle of this course.

---

## How The Course Uses This Separation

Every module from here on lives on one side of the divide or the other, or on the composition of the two:

| Module | Side |
|---|---|
| 2 — Parameterized Circuits | Objective (shape of the ansatz manifold) |
| 3 — Cost Landscapes | Objective (geometry) |
| 4 — Barren Plateaus | Objective (concentration phenomena) |
| 5 — Regularization | **Composition** — modifies dynamics in response to objective noise |
| 6 — Optimization Dynamics | Dynamics |
| 7 — Control-Theoretic View | Dynamics as controllers acting on objectives |
| 8 — Expressivity vs Trainability | Objective (a shape property, "expressivity", and a composition property, "trainability") |
| 9 — Hardware Noise Models | Objective (modifying the noise profile) |
| 10 — Adaptive Control Systems | **Composition** — closing additional feedback loops |

This split is not arbitrary. It follows from the claim of this node. If objective and dynamics were inseparable, each module would need to discuss both, and the course would be an unnavigable tangle. Because they are separable, you can study Module 3 without caring which optimizer will be used, study Module 6 without caring which Hamiltonian is being minimized, and then compose them in Module 5 and beyond.

---

## Connection to the Thesis

Module 5's central claim — regularization as dynamical modification, not prior injection — **depends entirely on this separation**.

The argument goes: prior injection is a statement about the objective (the regularizer rewrites $C$ to a new function $C + \lambda R$). Dynamical modification is a statement about the dynamics (the regularizer changes the update rule without changing $C$). These two operations have **different semantics** — they produce different fixed points, different trajectories, different convergence properties.

If objective and dynamics weren't separable, you couldn't distinguish them, and the thesis wouldn't have anything to argue against. The fact that they are separable means there is a real choice to be made about **where** a regularization technique operates, and that choice has consequences.

Node 5.5 makes this argument in full. What this node does is establish the **conceptual substrate** that makes it possible to state the argument at all.

---

## Structural Role

This is the **conceptual capstone** of Module 1. It is also the **organizing principle** of the remaining nine modules. Every time a later module draws a clean line between "landscape" and "optimizer," it is cashing in the separation established here.

If you are reviewing Module 1 and you remember only two things, remember:

1. **Node 1.1:** why minimizing $\langle\psi|H|\psi\rangle$ is a legitimate strategy at all.
2. **Node 1.9:** that the objective and the dynamics are two different things.

Everything else in the module is infrastructure for those two claims.

---

## Relations to Other Concepts

- **No free lunch theorems** — averaged over all possible objectives, all optimizers are equivalent. The separation this node argues for is a prerequisite for stating NFL at all.
- **Plant / controller decomposition** — classical control theory's analogous split. See Module 7.
- **Data / model separation in classical ML** — structurally similar (data defines the loss landscape; architecture + optimizer define the dynamics of fitting it).
- **Game theory** — payoff function (objective) vs strategy (dynamics).

---

## Worked Example — Swapping Optimizers on the Same Landscape

Take a fixed VQE instance: $H_2$ Hamiltonian, UCCSD ansatz, 14 parameters, 1024 shots per evaluation. Run three optimizers to 500 iterations each:

- **COBYLA** — gradient-free local search. Converges slowly but reliably; low seed variance.
- **Adam with parameter-shift gradients** — gradient-based. Converges faster per iteration but noisier; higher seed variance.
- **HOPSO** — population-based. Different exploration profile; best worst-case performance.

All three experience the **same objective**. Their different behaviors are entirely a property of the dynamics. If you aggregate the results as "VQE on $H_2$," you hide the fact that the variation across optimizers is larger than the variation across noise seeds — which means **the optimizer choice was the dominant variable**, not the landscape.

This is the empirical signature of Claim 2: dynamics matter, independent of the objective.

---

## Visualization

*Pending. Two-axis diagram. Horizontal axis: objective space (a cartoon landscape with hills and valleys). Vertical axis: trajectory space (curves through parameter space). Three different trajectories through the **same** landscape, one per optimizer. Annotation: "same objective, different dynamics, different outcomes." A companion panel: same optimizer on three different landscapes. Annotation: "same dynamics, different objectives, different outcomes."*

---

## Exercises

1. Construct a toy VQA example (2 parameters is enough) where two different optimizers starting from the same initialization reach different local minima. Identify whether the difference is attributable to the objective shape, the dynamics, or the composition.
2. Arrasmith et al. report that different optimizers produce qualitatively different convergence curves on identical VQE instances. Design a benchmark protocol that isolates the dynamics contribution from the objective contribution.
3. (VQA) The "barren plateau" phenomenon is typically described as a property of the objective. But for any fixed objective, some dynamics (e.g., very small step sizes) will still make progress while others (large step sizes) will not. Is "barren plateau" really a pure objective property, or is it a composition property? Defend your answer.

---

## Research Extensions

- **Objective-agnostic optimizer benchmarks** — protocols that report optimizer performance distributions across families of objectives, rather than single instances.
- **Dynamics-agnostic landscape analysis** — geometry, curvature, concentration results that hold regardless of which optimizer will be used.
- **Joint objective-dynamics co-design** — cases where the objective can be modified (via ansatz choice or measurement scheme) to improve the downstream dynamics. HOPSO thesis + regularization thesis together.

---

## Cross-links to Other Courses

- **Control Theory (Module 7 here)** — the plant-controller split.
- **Game Theory** — strategy vs payoff decomposition.
- **Systems Thinking** — the general pattern of "behavior as composition."
- **Statistics Module 10** — model selection as a search over objectives while holding dynamics fixed.

---

*Variational Quantum Algorithms · Module 1 — Hybrid Quantum Classical Systems · Node 1.9*
**Concepts:** cost-function, dynamical-systems, feedback
**Prerequisites:** 1.2, 1.5, 1.6, 1.8