

Either-And Power:
Designing Innovation With a Goal and a Plan

By Paul R. Leonard

The rapid expansion of computation is exposing an increasingly familiar constraint: electricity.

Artificial intelligence, data centers, advanced manufacturing, transportation, and electrification all increase demand for reliable electrical power. Communities and industries already connected to the grid experience the same constraint from a different direction. Additional generation requires capital, land, transmission, permitting, fuel, and time.

The natural response is to build more generation.

That is reasonable. It is also incomplete.

There is another question that can be asked without choosing between nuclear, natural gas, renewables, or any other generation technology:

Can we make existing energy infrastructure produce more useful electricity from the energy it already receives?

This is an either-and question.

It does not require replacing an existing generation technology. It does not require investors to abandon infrastructure already under development. It does not require choosing which primary energy source will dominate the future.

It asks whether an additional conversion mechanism can be placed within an existing energy system to recover energy that would otherwise leave that system as waste.

That distinction is important because innovation does not always require replacing what already works. Sometimes the more consequential innovation is the ability to add a new capability to an existing system without disturbing the system around it.

For thermal generation, this raises a particularly interesting possibility: direct solid-state generation of electricity from energy that is presently dissipated as heat.

The question is not whether heat engines work. They do.

The question is whether we have become so accustomed to converting energy through macroscopic thermal and mechanical cycles that we have overlooked another place to engineer the conversion: inside the material itself.

Optimization, Discovery, and Innovation

There is a useful distinction between optimization, discovery, and innovation.

Optimization and engineering operate within a system. They begin with a process, material, architecture, or technology and seek to make it work better. There is a plan, and the existing system provides the boundaries within which the plan operates. Materials are substituted, formulations refined, kinetics improved, variance reduced, manufacturing simplified, reliability increased.

Discovery expands knowledge. It does not require a goal, although it may involve an extraordinarily deliberate plan. A discovery can arise from a hypothesis, an experiment, an observation, or an unexpected result. Its immediate utility may be obvious, delayed, or entirely unknown.

Innovation requires both a goal and a plan.

The goal establishes something we want to accomplish that does not yet have an adequate solution. The plan provides a means of determining what must change to accomplish it.

The distinction matters because the three activities form a cycle rather than a hierarchy. Optimization can generate observations that become discoveries. Discovery can create new possibilities for innovation. And successful innovation eventually becomes optimization and engineering as the new system is made more reliable, efficient, economical, and useful.

The important point is that innovation creates something for optimization to work on.

There is another possibility that becomes especially interesting in materials science: the innovation may be the plan itself.

If a physical model can take a defined goal and environment of use and systematically construct candidate material architectures, while allowing unsuitable candidates to be rejected before extensive laboratory work, then the ability to make that plan is itself an innovation.

This is the direction in which Cascade's first-principles design methodology has developed.

The model does not eliminate laboratory work. It changes the role of the laboratory.

Early in the development of the methodology, more experimental work was required to answer questions concerning kinetics, particle sizes, interactions, synthesis behavior, and other variables that were difficult to establish beforehand. Each iteration provided information that tightened the model.

The initial design picture was useful but fuzzy around its edges.

With repeated use, those edges become more defined.

The practical objective is straightforward: answer more questions before entering the laboratory, and therefore spend less laboratory time determining which questions were worth asking.

That is an optimization of the innovation itself.

It is separate from optimizing the material or system eventually produced. Once a candidate architecture has been demonstrated, engineering can improve its formulation, manufacturing, durability, cost, operating range, and other characteristics within its environment of use.

The design methodology and the resulting technology therefore have different optimization cycles.

That distinction becomes useful when considering a new goal.

A Different Goal for the Same Design Method

Suppose the goal is not to design another catalyst.

Suppose the goal is:

Generate additional electricity directly from an existing thermal energy system without introducing another mechanical conversion cycle.

That is sufficiently specific to become an engineering problem, while broad enough to permit an unconventional solution.

The environment of use can be specified progressively.

An existing power-generating facility already receives a primary energy input. It converts part of that input into electricity and rejects another portion through thermal pathways. Those rejected flows eventually equilibrate with the surrounding environment.

The proposed innovation does not need to replace the primary generator.

It needs to interact with the energy that the existing system is already rejecting.

That changes the question from:

What new power plant should we build?

to:

What material architecture could be added to an existing power system to convert part of its otherwise rejected energy into additional electricity?

That is a very different design problem.

It also provides an unusually clean environment of use. The source, temperature range, physical interface, operating conditions, and desired electrical output can all be progressively specified. The problem can therefore be handed to a first-principles design methodology rather than beginning with a search through existing device categories.

The first question is then not which generator should we copy?

It is:

What must the material physically do to produce an extractable electromotive force under these conditions?

Start With the Desired Function

Electricity is electronic behavior.

The electromagnetic force governs the interactions responsible for electronic structure, and electronic behavior determines a remarkably broad range of material characteristics. Chemical reactivity, electrical and thermal conductivity, optical properties, magnetism, and mechanical bonding all depend, in different ways, upon how electrons occupy energy states, share or transfer valence electrons, and respond to their surrounding fields and lattice structure.

That makes electromagnetism a natural domain in which to begin a problem whose desired output is electrical.

This does not mean that any electronic excitation can become useful electricity. Nor does it imply that thermodynamics can be circumvented.

It means that the desired function can be expressed in physical terms.

A useful material architecture must receive energy from its environment, establish a non-equilibrium electronic condition, maintain some degree of separation or directional asymmetry, and provide a pathway through which that separation can become an extractable electrical potential.

The design problem therefore becomes one of architecture.

What electronic states are available?
How can the environment populate them?
How long do they persist?
How does energy move through the material?
Where can charge accumulate?
What prevents immediate recombination?
What structural features can create directional separation?
How can the resulting potential be coupled to an external circuit?

These are not questions that require beginning with a conventional generator and optimizing its components. They can be treated as material-design questions.

And this is where the distinction between macroscopic and microscopic energy gradients becomes important.

The Energy Gradient Inside the Material

Macroscopic thermal systems deal with energy distributed across enormous numbers of degrees of freedom. Their engineering architectures are consequently built around bulk temperature, pressure, fluid flow, phase changes, and mechanical work.

At the electronic scale, matter contains a different energy landscape.

The characteristic thermal energy at room temperature is approximately 0.025 eV. Electronic state differences can be substantially larger. The existence of these larger energy differences does not create energy from nothing; rather, it provides the architecture through which energy can be stored, transferred, and redistributed within matter.

The engineering opportunity is therefore not simply to create a larger temperature difference.

It is to investigate whether an energy input that would ordinarily become distributed thermal motion can instead be coupled, under appropriate material conditions, into persistent electronic non-equilibrium.

That is the relevant distinction in Cascade's work between ordinary macroscopic gradients and the much larger state-to-state energy differences available within an electronically structured material.

The question becomes one of lifetime and pathway.

If an excited electronic state immediately relaxes, its energy is rapidly redistributed into lower-energy modes and ultimately becomes heat. If the material architecture can alter the available relaxation pathways, however, the excitation may persist sufficiently long to participate in other processes.

This is where defects, interfaces, lattice structure, and electronic coupling become engineering variables rather than imperfections to be eliminated.

The objective is not maximum purity.

The objective is the right architecture for the desired function.

From Material Architecture to Charge Separation

Cascade's experimental record provides the foundation for asking this question.

The working material systems developed to date demonstrate that electronically complex, non-equilibrium behavior can be deliberately produced through material architecture and synthesis conditions. The relevant architectures are not dependent upon exotic elements or high-pressure fabrication. Cascade's demonstrated catalyst work has included systems in which readily available materials are used in place of elements traditionally associated with particular functions.

That history does not prove direct electricity generation.

It does something more useful for the present purpose: it establishes that there is a working design methodology from which the next question can be constructed.

For direct generation, the desired sequence would be conceptually simple:

energy input → electronic excitation → persistence and transport → directional charge separation → electromotive force → electrical output

The difficult part is making every step coexist within one material architecture.

The material must not merely absorb energy. It must manage what happens to that energy afterward.

A useful architecture could therefore involve a network of structural and electronic features that modifies the ordinary relaxation pathway. Defects and interfaces may provide intermediate states, trapping sites, coupling pathways, or barriers. Lattice vibrations can participate in energy transfer. Excited states can migrate through the material rather than immediately collapsing locally.

And if the architecture contains an appropriate asymmetry, charge separation need not remain purely local.

This is the conceptual basis for what Cascade has described as quantum catalytic separation of charges: using an engineered non-equilibrium material architecture to facilitate the separation of electrons and holes rather than allowing excitation to terminate immediately in recombination and heat.

The important word is engineered.

The desired charge separation is not being assumed to emerge simply because a material has defects or because it contains an excited state. The design problem is to determine what combination of electronic structure, defect architecture, spatial asymmetry, kinetics, and environment can produce a measurable and useful EMF.

That is exactly where a first-principles model has something to contribute.

The Material Does Not Need to Be Perfect

There is another reason this problem is particularly suited to Cascade's approach.

Conventional semiconductor engineering often treats structural precision and purity as central requirements because the desired electronic behavior depends upon tightly controlled interfaces and carefully defined material properties.

A different possibility exists when the desired function is produced by a self-organizing material architecture.

If the active state of a material can be produced through chemical self-assembly under appropriate synthesis conditions, then the architecture does not necessarily need to be constructed atom by atom or layer by layer. The material itself can find a stable configuration containing the structural and electronic features required for the function.

That possibility is significant because it changes the manufacturing problem.

The question becomes not:

How do we fabricate an impossibly precise quantum device?

but:

What synthesis conditions cause the material to assemble into the architecture that the design model requires?

This is consistent with the practical experience of Cascade's existing material development. Synthesis has been performed using standard laboratory equipment and atmospheric-pressure processing rather than requiring an exotic fabrication environment. The operating environment of the resulting systems has likewise generally involved atmospheric pressure and, depending upon the application, moderate temperatures.

That does not establish that every proposed direct-generation architecture will be equally simple to manufacture.

It establishes a useful design preference:

If the required electronic architecture can be produced through self-assembly, then manufacturing simplicity becomes part of the design problem rather than a downstream problem.

Chemical vapor deposition, thermal processing, precipitation, and other synthesis routes could become candidates depending upon the architecture ultimately identified. The appropriate method would be determined by the material and environment of use rather than selected in advance.

An Existing Asset as the Environment of Use

The commercial power of an *either-and* approach becomes apparent when considering the sheer volume of energy already moving through a working power facility. A modern thermal or nuclear asset is highly effective at its primary mission: converting a concentrated energy source into reliable electrical power. But that process also establishes substantial secondary energy flows through cooling systems, exhaust surfaces, heat exchangers, and other interfaces with the surrounding environment. These flows are normally treated as consequences of the primary conversion process and ultimately as part of its thermal rejection.

A different approach is to treat those flows as the **environment of use for a second energy-conversion architecture**.

This changes the engineering question. Conventional waste-heat recovery asks how efficiently rejected thermal energy can be passed through another macroscopic conversion cycle. Direct solid-state generation asks a more fundamental question: can the material architecture itself be designed to interact with the available energy at the electronic level, before that energy is fully degraded into the thermal reservoir?

In this framing, a portion of energy that has already become disordered thermal motion is used to establish a microscopic non-equilibrium electronic condition from which ordered electrical energy can be extracted.

Cascade's experimental work provides a basis for asking that question. We have observed non-equilibrium electronic behavior in material architectures in which defect structures, excited states, charge-transfer pathways, and interfaces are not incidental features of the material but functional components of its behavior. In these systems, the relevant energy landscape extends far beyond the approximately **0.025 eV thermal scale associated with room temperature**, with experimentally observed electronic energy differences reaching **5 eV and beyond**.

The significance is not that a 5 eV electronic state translates directly into 5 eV of electrical output. It does not. The significance is that the material architecture provides access to an energy scale vastly greater than the thermal fluctuation that initiates the process. The engineering problem therefore becomes one of capture, persistence, transport, and charge separation: how much of that electronically available energy can be retained within the architecture long enough to produce a useful and extractable electromotive force?

That is the basis for considering quantum catalytic separation of charges as a route to direct solid-state generation.

The proposition is deliberately different from simply adding another heat engine. The goal is to design a material subsystem in which thermal input initiates non-equilibrium electronic behavior, the architecture controls its subsequent evolution, and charge separation produces an electrical potential that can be coupled directly to an external circuit. The material is therefore not merely the working fluid of a conversion process. **The material is the conversion architecture.**

The scale of the opportunity follows from the environment of use. A working power facility already possesses the energy source, the physical infrastructure, the operating environment, and the electrical interconnection. If a solid-state architecture can be engineered to convert a fraction of an otherwise rejected energy flow directly into electricity, that output is additive to the facility's primary generation function. No replacement of the primary generating technology is required. No additional fuel is required for the secondary generation process. The existing asset simply acquires another productive function.

If the working model holds, a double-digit percentage increase in net electrical output is a reasonable engineering target for investigation. It should not be regarded as a claimed performance specification; it is a target derived from the scale of the underlying energy landscape and the quantity of energy already passing through the facility. And there is no reason at this stage to assume that double-digit improvement represents the upper bound. The purpose of the target is to establish a sufficiently consequential engineering problem to test.

This is where *either-and* becomes more than a rhetorical distinction.

The objective is not to determine whether nuclear, thermal, or another generation technology is the right answer. Those technologies already perform valuable functions at enormous scale. The objective is to determine whether the same physical assets can be made more productive by adding a fundamentally different conversion architecture to their environment of use.

The asset would not need to be replaced.

It would become more productive.

This Is Where the Innovation Loop Begins

At this point, the question becomes concrete enough to enter a Phase I feasibility process.

The goal is known.
The environment of use can be specified.
The required physical function can be described.
The existing experimental record provides relevant constraints and candidate mechanisms.

The first-principles model can then be used to construct and reject candidate architectures before committing substantial laboratory resources.

The initial output would not be a finished generator.

It would be a design picture.

Some portions of that picture may already be well defined. Others will remain fuzzy. The model should identify both.

That distinction matters because early falsification is more valuable than premature certainty.

If a proposed architecture cannot maintain the necessary non-equilibrium state, it can be rejected.
If charge separation cannot be sustained, it can be rejected.
If the expected energy coupling is insufficient under the actual environment of use, it can be rejected.
If synthesis cannot produce the required architecture reproducibly, that becomes a design constraint.

If the architecture survives those tests, it becomes an experimental candidate.

The laboratory then answers the questions that the model cannot yet answer with sufficient confidence.

And those answers return to the model.

This is the same cycle that has progressively tightened Cascade’s existing design methodology. The difference is that the goal has changed.

We are no longer asking the model to design a catalyst.

We are asking it to determine whether electricity itself can become the desired material function.

The Either-And System

If such an architecture proves viable, its significance does not depend upon replacing existing generation.

That is precisely the point.

A nuclear plant can remain a nuclear plant.
A gas turbine can remain a gas turbine.
An industrial process can remain an industrial process.

The innovation would add another conversion layer to the system.

Existing infrastructure would become the environment of use for a new class of material architecture.

This is economically important because technological transitions do not always need to destroy the capital that came before them. A technology that increases the productive output of an existing asset can be valuable precisely because it enhances rather than invalidates the investment already made.

The same logic applies at the national level.

If additional electricity can be produced from energy already moving through industrial infrastructure, the benefit is not limited to the generator itself. Additional electrical output can support additional computation, manufacturing, transportation, water treatment, communications, and local economic activity without requiring every increase in demand to be met by an entirely new primary generation facility.

The result is not necessarily a new energy system.

It is a more productive existing system.

That is the either-and.

We do not have to decide whether nuclear or another generation technology is the future before asking whether its rejected energy can become another source of useful output.

We can build generation and recover more of the energy flowing through it.
We can expand computation and increase the productivity of the infrastructure supplying it.
We can preserve large centralized generation assets and add new electrical functionality to their physical environment.

The innovation is additive.

The More Interesting Question

There is a temptation in advanced materials research to begin with the material.

What compound should we synthesize? What defect should we introduce? What interface should we fabricate? What property can we improve?

A goal-oriented first-principles approach reverses that sequence.

Start with the desired function.
Define the environment in which it must operate.
Determine the physical behavior required.
Then ask what architecture can produce that behavior.

For direct solid-state electricity generation, the goal is sufficiently clear to make this process meaningful. The proposed mechanism is sufficiently grounded in Cascade's existing experimental record to make the questions worth asking. And the unknowns are sufficiently identifiable that the next steps can be framed as falsifiable engineering questions rather than open-ended exploration.

That is the real opportunity.

The objective is not to claim that the answer has already been found.

It is to demonstrate that the problem can be planned.

A first-principles design model becomes more valuable as it answers more questions before the laboratory is involved. The innovation is therefore not a static piece of intellectual property. It is a capability that improves through use.

Today, portions of the design space remain fuzzy.

Tomorrow, fewer will.

And if that progression continues, the distinction between "researching what might work" and "engineering what will work" becomes increasingly consequential.

The first goal of this exercise, therefore, is not to build a new power plant.

It is much simpler:

Take an energy system that already exists, identify energy that it currently rejects, and determine whether a material architecture can be designed to convert some of that energy directly into useful electricity.

That is a sufficiently bounded problem to plan.
It is sufficiently physical to falsify.
And, if successful, it becomes something very useful indeed:

a new optimization problem waiting to begin.

#EitherAndPower #AdditiveGeneration #MaterialsInnovation #DirectElectricity #EnergyInfrastructure

###

References:

[1] Kopylov, V. B.; Leonard, P. R. *Emergent Materials Design and Electronically Excited-State Catalysis.* Cascade Catalysts, Inc.
DOI: https://zenodo.org/records/18115403

[2] Kopylov, V. B.; Leonard, P. R. *Can Smart Matter Be Modeled by Design? Turning Passive Solids into Active Systems.* Cascade Catalysts, Inc.
DOI: https://zenodo.org/records/1779153

Paul R. Leonard writes Big Quantum Nerd, a LinkedIn blog exploring structural clarity through a thermoeconomic lens. His work traces seemingly disparate issues — homelessness, innovation, affordability, healthcare, sovereignty, and institutional failure — back to their common root in energy pathways and surplus dynamics, and how non-equilibrium design principles can restore human agency and resilience.

Paul is co-founder, President & CEO of Cascade Catalysts, Inc., a company developing quantum-effect catalysts and advanced materials for industrial applications. He is also founder of Activate Biophysics Corp., advancing non-toxic, systems-level health technologies. Drawing on decades of applied work across physics, energy, materials, and complex systems, Paul works with leaders in government, industry, and research to move beyond symptom management toward structural solutions that expand surplus and amplify human effort rather than consume it.