

# Getting to Grips with Sweedler Runners for Systems

Appendix B: more about good regulators

September 2026

## 1 Good regulators, internal models, and mutual modelling

In Parts III and V we considered the good-regulator theme through the Moore–Mealy construction of Virgo et al. There the basic picture was a closed-loop system obtained by wiring together a Moore machine and a Mealy machine. At the level of theories, the interaction was organised by a map of the form

$$T_{\text{tick}} \longrightarrow T_{\text{Moore}} + T_{\text{Mealy}},$$

so that a common ticking process drove both components. The Moore and Mealy machines themselves were not merely theories. They were stateful realizations of those theories: comodels, or more generally runners, carried by state spaces. This distinction will be important throughout the present appendix. Passing to a predicate-enriched setting then allowed us to ask whether the resulting closed-loop stateful system preserved a desired condition. In the simplest case, this was a safety predicate on the carrier of the closed system. That construction already captured an important part of the good-regulator story: a regulator is not studied in isolation, but as part of a wired system whose behaviour can be constrained by predicates. We can now revisit the same idea with some of the extra language developed since then. The main refinement is to distinguish three things which were previously compressed into the one regulator component:

1. a stateful realization of the system being regulated;
2. a stateful internal model carried by the regulator;
3. the policy or mechanism by which the regulator acts on the system.

This gives a more direct way to formulate the claim that a good regulator must in some sense model the system it regulates.

### 1.1 Plant, model and regulator

Let

$$T_P$$

be a theory describing the relevant operations of the system to be regulated. The plant itself is a stateful realization of this theory. In the simplest case it is a comodel

$$T_P \rightsquigarrow \text{Id}$$

carried by a set

$$S_P.$$

More generally, if the plant runs against some background theory  $R_P$ , it may be represented by a runner

$$T_P \rightsquigarrow R_P$$

with residual state  $S_P$ . Thus  $S_P$  is not a state space belonging to the theory  $T_P$  itself. It is the carrier or residual state of a particular realization of that theory. We shall call this stateful realization the *plant*. Likewise, let

$$T_{\hat{P}}$$

be the theory used by the regulator for its internal representation of the plant. The internal model is again not the theory alone. It is a stateful realization of that theory, carried by a set

$$S_{\hat{P}}.$$

For example, it may be a comodel

$$T_{\hat{P}} \rightsquigarrow \text{Id}$$

or a more general runner with residual state  $S_{\hat{P}}$ . Finally, let

$$T_R$$

describe the regulator's own control operations, again realized by an appropriate stateful system. In the earlier Moore–Mealy example, the regulator was represented directly by a Mealy comodel. We now allow that regulator to carry, as part of its own stateful realization, an additional model state

$$\hat{s} \in S_{\hat{P}}.$$

The first question is therefore not yet whether the regulator succeeds, but whether its internal stateful model is adequate to the actual plant.

## 1.2 Agreement on observations

The weakest requirement is that the plant realization and the model realization agree on the observations relevant to regulation. Suppose both can be observed through some common observation theory

$$T_O.$$

Concretely, the plant comodel on  $S_P$  and the model comodel on  $S_{\hat{P}}$  may each induce an observation map

$$\text{obs}_P : S_P \longrightarrow O,$$

and

$$\text{obs}_{\hat{P}} : S_{\hat{P}} \longrightarrow O.$$

If

$$s \in S_P, \quad \hat{s} \in S_{\hat{P}},$$

we can ask that

$$\text{obs}_P(s) = \text{obs}_{\hat{P}}(\hat{s}).$$

This already gives a useful version of the good-regulator principle:

the regulator need not reproduce the whole plant state;  
it must reproduce the distinctions relevant to regulation.

The internal model realization may therefore have a much smaller carrier than the plant realization. This is important. A good regulator need not contain an exact copy of the system it regulates. It needs a stateful model which is adequate for the task.

### 1.3 A simulation predicate

A stronger formulation compares the carrier of the plant realization with the carrier of the model realization directly. Let

$$R \subseteq S_P \times S_{\hat{P}}$$

be a relation saying that the model state  $\hat{s}$  is an adequate representation of the plant state  $s$ . For example,

$$R(s, \hat{s})$$

might mean that the two realizations agree on a finite collection of relevant observables, or that their outputs differ by no more than some permitted tolerance. Passing to a predicate-enriched setting, we can require this relation to be preserved by the coupled plant/model dynamics:

$$R(s, \hat{s}) \implies R(s', \hat{s}').$$

Thus the regulator carries not merely a state estimate, but a stateful realization whose carrier remains related to the carrier of the actual plant. This gives a useful refinement of the construction from Part V. There we used predicates chiefly to express desirable behaviour of the closed-loop system. Here we use a predicate one level earlier, to express the adequacy of the model on which regulation depends. We therefore obtain two kinds of invariant:

$$\boxed{R(s, \hat{s}) \quad \text{and} \quad \text{Safe}(s).}$$

The first says that the model realization remains good enough. The second says that the plant realization remains within the desired region. A regulator may therefore fail in two quite different ways:

- the model state may cease to track what matters about the plant state;
- the model may remain adequate, but the resulting control policy may still fail to maintain the desired predicate.

Separating these two failures is one advantage of the richer formulation.

### 1.4 The internal model as a runner

There is also a more explicitly Sweedler-style way to express the internal model. Suppose the regulator does not merely carry a comodel of some abstract model theory, but a stateful interpretation of the plant theory itself. Then one may ask for a runner

$$T_P \rightsquigarrow T_{\hat{P}}.$$

If its residual state is  $S_{\hat{P}}$ , this is represented by a coalgebra

$$S_{\hat{P}} \longrightarrow \{T_P, T_{\hat{P}}\}(S_{\hat{P}}).$$


Operationally, this says:

the regulator carries a stateful procedure for interpreting plant operations inside its own model theory.

The state space  $S_{\hat{P}}$  is therefore not an additional state space attached to the theory  $T_{\hat{P}}$ . It is the residual state of the runner which realizes the modelling relation. This makes the phrase “the regulator contains a model of the plant” rather literal. A plant operation can be run internally against the model state, producing a corresponding model computation and an updated model state. The model need not reproduce every operation of the plant. In applications one would normally choose a theory  $T_P$  containing only those operations relevant to the regulation problem, or else pass to an appropriate quotient or observation theory. So the interesting question is not whether

$$T_P \rightsquigarrow T_{\hat{P}}$$

is an exact representation, but whether the resulting runner preserves enough structure for the desired predicates. This suggests the hierarchy



The good-regulator claim concerns the passage through this hierarchy: successful regulation requires a stateful model which preserves the distinctions needed by the control task.

### 1.5 A minimal temperature example

A tiny example makes the distinction concrete. Let the plant be a comodel carried by a temperature state

$$T \in S_P = \mathbb{Z}.$$

Let the regulator carry an internal model runner or comodel with carrier

$$S_{\hat{P}} = \mathbb{Z},$$

whose current model state we write as

$$\hat{T}.$$

Suppose the desired safety predicate on the plant carrier is

$$18 \leq T \leq 22.$$

The regulator does not need exact equality of the two states. Perhaps it is sufficient that

$$|T - \hat{T}| \leq 1.$$

Define

$$R(T, \hat{T}) \quad :\Longleftrightarrow \quad |T - \hat{T}| \leq 1.$$

Then two questions are clearly separated:

|                 |                                   |
|-----------------|-----------------------------------|
| model fidelity: | $R(T, \hat{T})$ remains true;     |
| regulation:     | $18 \leq T \leq 22$ remains true. |

The first concerns the relation between two stateful realizations. The second concerns the relation between the plant realization and its specification. A good closed-loop design must preserve both. This is only a toy example, but it already shows why the internal-model layer is useful. Without it, a single closed-loop safety predicate tells us whether regulation succeeds, but not whether it succeeds because the regulator carries an adequate stateful model of the plant.

### 1.6 From internal models to IntRun

The picture becomes more interesting when the regulator does not merely run a private simulation and then issue an independent command. Suppose plant behaviour is processed through both an internal-model runner and a control runner. Schematically, one may encounter a prism-like arrangement in which plant behaviour is interpreted using

$$T_{\hat{P}} + T_R.$$

The internal-model runner provides a stateful interpretation of what the plant is doing; the control component provides the language in which the regulator acts back on the plant. In an IntRun-style picture these roles can belong to one stateful wiring component. One face carries the model theory, another carries the regulator theory, and the plant-facing theory is handled through their combined behaviour. The precise placement of the theories depends on which direction of interaction we wish to emphasise, but the operational picture is simple:

plant behaviour enters the regulator; it is interpreted by a stateful internal model; the resulting model state contributes to the control behaviour sent back to the plant.

The important point is that the model state belongs to a runner or comodel inside the wiring. It is not an extra state space floating beside the formal construction. This begins to look much closer to predictive control than the simple Moore–Mealy loop considered earlier.

### 1.7 Handlers and regulation

The handler viewpoint gives another reading. Suppose the plant presents operations in a theory  $T_P$ , while the regulator’s internal resources provide handlers for the model theory and the control theory. An IntRun component can combine these handlers to produce an effective handler for plant behaviour. Thus regulation can be understood as a special case of the general picture from Appendix A:

|                                                                                      |
|--------------------------------------------------------------------------------------|
| component + available handlers $\longrightarrow$ new handler for the exposed theory. |
|--------------------------------------------------------------------------------------|

In the present case, the internal-model runner contributes to the handler by deciding how observed plant behaviour should be interpreted before the control response is chosen. This makes the good-regulator condition a condition not merely on a state estimate, but on the quality of the induced handler. Two internal-model runners may differ substantially while inducing the same behaviour on all plant operations relevant to the regulation task. From the point of view of regulation, they may then be equally good. Again, the appropriate notion of modelling is task-relative.

## 1.8 Mutual modelling

The same architecture becomes more interesting when both systems are active agents. Suppose agents  $A$  and  $B$  are represented by stateful realizations of theories

$$T_A, \quad T_B,$$

carried respectively by

$$S_A, \quad S_B.$$

Agent  $A$  may carry an internal model of  $B$  in the form of a runner

$$T_B \rightsquigarrow T_{\hat{B}}^{(A)}$$

with residual state

$$\hat{S}_B^{(A)}.$$

Likewise,  $B$  may carry a model of  $A$ :

$$T_A \rightsquigarrow T_{\hat{A}}^{(B)}$$

with residual state

$$\hat{S}_A^{(B)}.$$

We then have two model-adequacy predicates,

$$R_A(s_B, \hat{s}_B), \quad R_B(s_A, \hat{s}_A).$$

The behaviour of each agent depends not only on the stateful realization of the other but on the internal runner through which that other is interpreted. Thus the coupled system has, at minimum, carrier

$$S_A \times S_B \times \hat{S}_B^{(A)} \times \hat{S}_A^{(B)}.$$

This creates a phenomenon absent from the one-way regulator example. An internal model can help to produce the behaviour which appears to confirm that same model. For instance,  $A$  may expect a certain response from  $B$ , act on the basis of that expectation, and thereby elicit exactly the response predicted. The model may therefore be dynamically stable without being especially accurate outside the interaction it helps create. This suggests that in mutually regulating systems we should distinguish

predictive success

from

independent adequacy.

A model may predict well because it correctly represents the other stateful system, or because the surrounding interaction repeatedly drives the other system into the region the model already expects. The predicate-enriched language gives us a place to formulate such distinctions.

## 1.9 A relational residual state

As in Appendix A, it may be useful to introduce an additional residual state

$$S_{AB}$$

belonging to the interaction rather than to either agent alone. This can arise as the residual state of a runner mediating the interaction between the two systems. It can record the accumulated history of the interaction: expectations, commitments, unresolved episodes, conventions, or simply whatever information is needed for the coupled running process. Then the full carrier becomes

$$S_A \times S_B \times S_{AB} \times \hat{S}_B^{(A)} \times \hat{S}_A^{(B)}.$$

A recurrent interactional pattern can now be understood as a recurrent region or invariant of this larger state space. In particular, one can ask whether the mutual-model predicates

$$R_A, \quad R_B$$

are preserved, whether a relational predicate

$$R_{AB}$$

is preserved, and whether these in turn imply some desired system-level property. Thus the good-regulator question generalises naturally:

when one stateful system acts on another through an internal runner,  
what must that runner preserve for the desired interactional predicate to remain true?

## 1.10 Back to the good-regulator theorem


Seen from this perspective, the Moore–Mealy construction from Parts III and V is the simplest member of a larger family. At the first level we have

$$\text{plant comodel} + \text{regulator comodel} \longrightarrow \text{closed-loop comodel}.$$

At the next level we distinguish

$$\text{plant realization} + \text{internal-model runner} + \text{controller realization}.$$

At the next we allow both sides to carry stateful runners modelling one another. And in each case predicates can be used to express the adequacy conditions that must be preserved. The general pattern is therefore



This gives a natural Sweedler-style reading of the good-regulator principle. A regulator does not need to reproduce the entire system it regulates. What it needs is a stateful model—represented by a suitable comodel or runner—whose induced running behaviour preserves enough of the regulated system to support the predicates relevant to the control task. In this form, the good-regulator theme is no longer confined to classical control examples. It applies equally naturally to adaptive systems, predictive controllers, interacting agents, and other situations in which one stateful system must act successfully on another through an internal representation of its behaviour.