Continuation Value Is All You Need:

Reference Description

Jeffrey Sun

2026-05-17

Summary

Neural Value Function Iteration (nVFI) is an algorithm for solving value functions in heterogeneous-agent models with aggregate shocks. It consists of the following steps:

1. Parameterize the end-of-period value function (a function of both idiosyncratic and aggregate state) using a flexible nonlinear V estimator. (Reference implementation: a neural network.)
2. Draw samples from the aggregate state. (Reference implementation: from the ergodic distribution induced by the current V estimator.)
3. For each aggregate sample, compute the outer Bellman residual—a measure of the failure of the Bellman equation to hold at that aggregate state—using lookahead.
4. Update the V estimator to reduce the outer Bellman residuals across the sampled aggregate states. (Reference implementation: gradient descent on the parameters of the V estimator neural network.)
5. Evaluate global convergence of the V estimator. Return to Step 2 unless the criterion is met.

The algorithm is based on the following insight: if we choose a particular timing convention, then we can define a particular value function object $\mathcal{V}$, the **continuation value operator**.

I then show that: if we know $\mathcal{V}$, then we can solve the rest of the model using conventional methods.

This is directly analogous to how, in standard dynamic programming, knowing the value function allows us to solve the model. My view is that $\mathcal{V}$ is the natural generalization of the standard value function to heterogeneous-agent models with aggregate shocks. This is the meaning of the title, "Continuation Value Is All You Need".

1 Setup

1.1 Model Class and Timing Conventions

In this section, I specify the class of models for which the algorithm is defined. I specify a timing convention in order to cleanly separate out **aggregate shocks** from the **otherwise-deterministic** evolution of the aggregate state.

The setup is fairly general for discrete-time models, although it may require some translation (without loss of generality) for models written down with different timing conventions.

I also assume throughout that the model is stationary, ergodic, and has a unique global solution.

Time and the household period. Time is discrete and indexed by t. Each period consists of a *household period*, and a moment between household periods. Aggregate shocks are allowed to occur in this moment between household periods.[1]

It may occasionally be useful to refer to intermediate value functions or aggregate states within the period, which will also be denoted with a superscript, e.g. $V^{\text{postincome}}$. These are formally well-defined in terms of stage-decomposed household blocks (see Stages paper).

State. At the start of the period, each household has idiosyncratic state $x^{\text{start}} \in X^{\text{start}}$. At the end of the period, each household has idiosyncratic state $x^{\text{end}} \in X^{\text{end}}$.[2] The aggregate state at the start of period t is Γ_t^{start}; at the end, Γ_t^{end}.

[1] Henceforth, I may refer to the beginning (end) of the household period as the "start" ("end") of the period, with the aggregate shock taken to occur "between" periods.

[2] In Krusell–Smith, $x^{\text{start}} = (k^{\text{start}}, z^{\text{start}})$ and $x^{\text{end}} = (k^{\text{end}}, z^{\text{end}})$.

The aggregate state Γ_t^{start} consists of the cross-sectional household state distribution Λ^{start} over x^{start} and other aggregate state variables S^{start}, (and similarly for Γ_t^{end}),

$$\Gamma_t^{\text{start}} = (\Lambda_t^{\text{start}}, S_t^{\text{start}}) \tag{1}$$

$$\Gamma_t^{\text{end}} = (\Lambda_t^{\text{end}}, S_t^{\text{end}}). \tag{2}$$

Aggregate Shocks. An **aggregate shock** is defined as any shock to the model that:

1. Affects the aggregate state of the model, and
2. Is not perfectly foreseen.

Examples include:

- TFP shocks.
- Information revelation about future TFP, in a model with forward-looking behaviour, even if current TFP is unaffected.[3]
- Idiosyncratic productivity shocks to agents who are *granular* as opposed to *atomistic*. That is, shocks to agents whose individual behaviour can affect the aggregate state of the model.

The **only way** aggregate shocks are permitted to enter in this setup is:

1. As unforeseen changes in S_{t+1}^{start} relative to S_t^{end},
2. Occurring between the end of period t and the start of period $t+1$.

While this may seem restrictive, it is fairly general:

- If aggregate shocks occur in the middle of the period in a model, since time is periodic, the period boundary can simply be redefined such that the "period" ends immediately before the shock and begins immediately after.
- If information regarding a shock is revealed n periods in advance, this can be formulated as a shock to an information set which is included in the $S^{\cdot}$.

[3]This is an aggregate shock because it may change behaviour in the present.

- If shocks "immediately" affecting agents is part of the model, these can be formulated as shocks to the information set contained in S^{start}, which are actually physically resolved only within the period.

I also impose that, for any individual household, $x_t^{\text{end}} = x_{t+1}^{\text{start}}$ through any $t \to t+1$ transition.

Aggregate law of motion between periods. Given these timing assumptions for the aggregate shock(s), we can write the law of motion for S_t without loss of generality as,

$$\Gamma_{t+1}^{\text{start}} = (\Lambda_{t+1}^{\text{start}}, S_{t+1}^{\text{start}}) \tag{3}$$

$$\text{where} \quad \Lambda_{t+1}^{\text{start}} = \Lambda_t^{\text{end}} \tag{4}$$

$$S_{t+1}^{\text{start}} = \Omega(S_t^{\text{end}}, Z_t) \tag{5}$$

$$Z_t \sim \mathcal{D}, \tag{6}$$

where Z_t is the "random part" or innovation in the $S_t \to S_{t+1}$ transition, and $\mathcal{D}$ is its distribution.

In one line,

$$\Gamma_{t+1}^{\text{start}} = (\Lambda_t^{\text{end}}, \Omega(S_t^{\text{end}}, Z_t)), \quad \text{where} \quad Z_t \sim \mathcal{D}. \tag{7}$$

Information revelation. This setup implies the following fact, which we will use: The evolution of the aggregate state *within* the household period, from Γ_t^{start} to Γ_t^{end}, is **deterministic**.

In other words, the aggregate state at the end of the period, Γ_t^{end}, is uniquely determined by the aggregate state at the start of the period, Γ_t^{start}.[4]

2 Value Functions

We have multiple notions of "the value function" to deal with. In this section I define and disambiguate them.

[4]The endogenous determination of equilibrium quantities such as prices, quantities, populations, market tightnesses, entry, etc., are simply part of the process by which Γ_t^{end} is determined from Γ_t^{start}. In particular, while Γ_t^{end} may *depend* on prices in the sense of behaviour depending on prices, prices themselves are not a state variable contained in Γ^{start} or Γ^{end}.

Full global value function. The full global value function at the start of the period, V^{start}, is a function of both the idiosyncratic state x^{start} and the aggregate state Γ^{start},[5]

$$V^{\text{start}} : (x^{\text{start}}, \Gamma^{\text{start}}) \mapsto v \in \mathbb{R}.$$

Continuation value. We now break with common convention in an **extremely** useful way.

Recall that Γ^{end} is uniquely determined by Γ^{start} by our construction. Because of this, instead of writing the end-of-period value V^{end} in terms of Γ^{end}, as is common, we can write it in terms of Γ^{start} instead.[6] Absusing notation slightly,[7]

$$V^{\text{end}}(x^{\text{end}}, \Gamma^{\text{start}}) \equiv V^{\text{end (common convention)}}(x^{\text{end}}, \Gamma^{\text{end}}(\Gamma^{\text{start}})). \tag{8}$$

While this may seem misaligned, the story is intuitive. A household who begins the period in state x^{start} can only directly observe x^{start} and Γ^{start}, but must make decisions which depend on x^{end}-contingent values at the **end** of the period. $V^{\text{end}}(x^{\text{end}}, \Gamma^{\text{start}})$ intuitively reflects the act of forecasting end-of-period contingent values from the perspective of the start of period t.[8]

My entire strategy is based on this **continuation value** formulation. The core insight is that, if we have access to $V^{\text{end}}(\cdot, \Gamma^{\text{start}})$, or an approximation thereof, then the rest of the problem can be solved using conventional methods.

At this point, we can introduce the **continuation value operator** $\mathcal{V}$,

$$\mathcal{V}\Gamma^{\text{start}}(x^{\text{end}}) \equiv V^{\text{end}}(x^{\text{end}}, \Gamma^{\text{start}}). \tag{9}$$

My perspective is that this $\mathcal{V}$ is the natural generalization of the standard value function to heterogeneous-agent models with aggregate shocks.

[5] Explicitly, the expected NPV of future utility for a household in state x^{start} at the start of the period, given the aggregate state Γ^{start}.

[6] Why can we do this? V^{end} is obviously uniquely determined by x^{end} and Γ^{end}. But Γ^{end} is uniquely determined by Γ^{start} by construction (see previous section). Therefore, V^{end} is uniquely determined by x^{end} and Γ^{start}.

[7] Explicitly, $V^{\text{end}}(x^{\text{end}}, \Gamma^{\text{start}})$ refers to the NPV of future utility for a household in state x^{end} at the end of the period, given the aggregate state Γ^{start} at the start of the period.

[8] For example, a model in which households make relatively simple (e.g. consumption-savings) decisions might be written as,

$$V^{\text{start}}(x^{\text{start}}, \Gamma^{\text{start}}) = \max_a u(a; x^{\text{start}}) + \beta V^{\text{end}}(x^{\text{end}}(a; x^{\text{start}}), \Gamma^{\text{start}}).$$

Inner value functions. Note that objects like $\mathcal{V}\Gamma^{\text{start}}$ is a function of x^{end} only.[9] I will call such objects "inner value functions" in order to distinguish them from global value functions which also depend on Γ^{start}.

In the following, I will denote such inner value functions, that is any valuation over the state space X^{end}, by $\widetilde{V}^{\text{end}}$,[10]

$$\widetilde{V}^{\text{end}} : x^{\text{end}} \to \mathbb{R}.$$

It may help to think of $\widetilde{V}^{\text{end}}$ as a vector over a discretized state space X^{end}. This is how it is often represented on the computer: a vector of values at the gridpoints of x^{end}.

Value function approximations. The curse of dimensionality bites because it is impractical to represent V^{end} as a vector over an unavoidably large discretized state space $X^{\text{end}} \times \mathbb{R}^{\Gamma^{\text{end}}}$. Our strategy is to use a parameterized function approximator to approximate V^{end} instead. Given a guess of parameters θ, we denote the parameterized approximation by,

$$\widehat{V}^{\text{end}}_{\theta}(x^{\text{end}}, \Gamma^{\text{start}})$$

and the parameterized approximation of the continuation value operator by $\mathcal{V}_{\theta}$,

$$\mathcal{V}_{\theta}\Gamma^{\text{start}}(x^{\text{end}}) \equiv \widehat{V}^{\text{end}}_{\theta}(x^{\text{end}}, \Gamma^{\text{start}}).$$

In my reference implementation, $\mathcal{V}_{\theta}$ is a neural network taking a vector Γ^{start} to a vector $\widetilde{V}^{\text{end}}$, where the vectors are over discretized state spaces X^{start} and X^{end}.[11]

3 Continuation Value is All You Need

In this section I state my central claim: If $\mathcal{V}$ is known, then the rest of the model can be solved with conventional methods.

[9] In computer science, this is sometimes call a "partially applied function." The input Γ^{start} has been applied, but not the input x^{end}.

[10] This can be any function from the state space x^{end} to the reals. We will typically interpret this object as a guess of the slice of V^{end} at a given Γ^{start}, but note that it **does not make sense** to refer to a "true" value of $\widetilde{V}^{\text{end}}$ since $\widetilde{V}^{\text{end}}$ is ultimately still a variable.

[11] Note that this is also implicitly done by Krusell and Smith (1998), with the differences being that they use a linear interpolation over a grid of X^{end} and moments of Γ^{start}, and train the parameters θ indirectly via fitting laws of motion for aggregates.

3.1 Setup and Assumptions

An important assumption we need is that aggregate uncertainty is the reason why the model is intractable. That is, in the absence of aggregate uncertainty, the model would be tractable.[12] Explicitly,

Assumption 1. *If we shortened the time horizon of the model to one period, with exogenously-given initial aggregate state* Γ^{start} *and exogenously-given end-of-period* x^{end}*-contingent payoffs* $\widetilde{V}^{end}(x^{\text{end}})\ \forall x^{\text{end}}$*, then we could use conventional methods to find all necessary equilibrium objects (e.g. prices, quantities), obtain* $\widetilde{V}^{start}(x^{\text{end}})$ *by backwards iteration, and simulate forward to obtain* Γ^{end}*.*

Assumption 1 can be rephrased as assuming that we have access to a **within-period solver** Φ which takes a start-of-period aggregate state Γ^{start} and an end-of-period inner value function $\widetilde{V}^{\text{end}}$, and returns an end-of-period aggregate state Γ^{end} and start-of-period inner value function $\widetilde{V}^{\text{start}}$,

$$\Phi:\ (\Gamma^{\text{start}},\ \widetilde{V}^{\text{end}})\ \longmapsto\ (\Gamma^{\text{end}},\ \widetilde{V}^{\text{start}}).$$

For convenience, let us also define,

$$\Phi_{\Gamma}(\Gamma^{\text{start}},\ \widetilde{V}^{\text{end}}) \equiv \Phi(\Gamma^{\text{start}},\ \widetilde{V}^{\text{end}})_1 = \Gamma^{\text{end}}$$
$$\Phi_{V}(\Gamma^{\text{start}},\ \widetilde{V}^{\text{end}}) \equiv \Phi(\Gamma^{\text{start}},\ \widetilde{V}^{\text{end}})_2 = \widetilde{V}^{\text{start}}.$$

3.2 Within-Period Solution Given $\mathcal{V}$

Proposition 1. *If we have access to the true* $\mathcal{V}$*, and Assumption 1 holds, then the model can be solved. In particular, the model can be simulated forward correctly and value functions can be computed for any state* Γ^{start}*.*

Proof. Given Γ^{start}, we can obtain $V^{\text{end}}(\cdot, \Gamma^{\text{start}}) = \mathcal{V}\Gamma^{\text{start}}$.

The within-period solver yields,

$$\Phi(\Gamma^{\text{start}}, V^{\text{end}}(\cdot, \Gamma^{\text{start}})) = (\Gamma^{\text{end}}, V^{\text{start}}(\cdot, \Gamma^{\text{start}})).$$

[12] If the model is intractable even in the absence of aggregate uncertainty, then the issue lies elsewhere and the algorithm described here will not be helpful.

Finally, by Equation 7, for any Z, we can simulate forward to the beginning of the following period,

$$\Gamma^{\text{start}'} = (\Lambda^{\text{end}}(\Gamma^{\text{end}}), \Omega(S^{\text{end}}(\Gamma^{\text{end}}), Z)).$$

Thus, we are able to simulate the model forward by one full period, and can therefore simulate forward by arbitrarily many periods. □

4 Outer Bellman Equation and Operator

Outer Bellman Equation. A model in our setup can thus be abstractly described by the following four equations,

$$\forall x^{\text{end}}, \Gamma^{\text{start}}, \qquad \mathcal{V}\Gamma^{\text{start}}(x^{\text{end}}) = \mathbb{E}_Z[V^{\text{start}'}(x^{\text{end}}, \Gamma^{\text{start}'})] \tag{10}$$

$$\text{where} \quad V^{\text{start}'}(\cdot, \Gamma^{\text{start}'}) \equiv \Phi_V(\Gamma^{\text{start}'}, \mathcal{V}\Gamma^{\text{start}'}) \tag{11}$$

$$\Gamma^{\text{start}'} \equiv \Omega(\Gamma^{\text{end}}, Z) \tag{12}$$

$$\Gamma^{\text{end}} \equiv \Phi_\Gamma(\Gamma^{\text{start}}, \mathcal{V}\Gamma^{\text{start}}). \tag{13}$$

Equation 10 can be viewed as the **inner Bellman Equation**. It defines V^{end} as the expectation over realizations of the aggregate shock Z of the start-of-next-period value $V^{\text{start}'}$.

Equation 11 defines start-of-next-period $V^{\text{start}'}$ as the result of obtaining end-of-next-period $V^{\text{end}'}$ via $\mathcal{V}$, then backwards-iterating within-period via Φ_V.

Equation 12 defines $\Gamma^{\text{start}'}$ as the result of simulating Γ^{end} through the stochastic transition Ω.

Equation 13 defines Γ^{end} as the result of simulating Γ^{start} forward to the end of the period, given V^{end}.

The four equations taking jointly can be viewed as an **outer Bellman Equation** implicitly defining the **continuation value operator** $\mathcal{V}$.

Outer Bellman Operator. The left hand side of Equation 10 can be thought of as forecasting the end-of-period-t value function from the beginning of period t via $\mathcal{V}$.

The right hand side of Equation 10, together with equations 11-13, can be thought of as forecasting the end-of-period-t value function from the beginning of period t via forecasting to

the end of period $t+1$, then backwards-iterating to the end of period t.

Together, they state that these two operations must give the same result.

Explicitly, we can define an **outer Bellman Operator**,[13]

$$\begin{aligned}
(\mathcal{L}\mathcal{V})\Gamma^{\text{start}}(x^{\text{end}}) &\equiv \mathbb{E}_Z[V^{\text{start}'}(x^{\text{end}}, \Gamma^{\text{start}'})] \\
\text{where} \quad V^{\text{start}'}(\cdot, \Gamma^{\text{start}'}) &\equiv \Phi_V(\Gamma^{\text{start}'}, \mathcal{V}\Gamma^{\text{start}'}) \\
\Gamma^{\text{start}'} &\equiv \Omega(\Gamma^{\text{end}}, Z) \\
\Gamma^{\text{end}} &\equiv \Phi_\Gamma(\Gamma^{\text{start}}, \mathcal{V}\Gamma^{\text{start}}).
\end{aligned}$$

The **outer Bellman Equation** can then be written,

$$\mathcal{V} = \mathcal{L}\mathcal{V}. \tag{14}$$

The Neural Bellman Objective What I call **Neural Value Function Iteration** (nVFI) simply refers to training a neural network to satisfy Equation 14. Explicitly,

$$\min_\theta \mathbb{E}_{\Gamma^{\text{start}}, x^{\text{end}}}\left[(\mathcal{V}_\theta\Gamma^{\text{start}}(x^{\text{end}}) - \mathcal{L}\mathcal{V}_\theta\Gamma^{\text{start}}(x^{\text{end}}))^2\right]. \tag{15}$$

5 nVFI Algorithm

Step 1 — Parameterize $\mathcal{V}$ and guess θ.

Choose a flexible nonlinear estimator $\mathcal{V}_\theta$ that takes an aggregate state Γ^{start} as and returns a vector (or other object) approximating $V^{\text{end}}(\cdot, \Gamma^{\text{start}})$. Initialize θ.

Step 2 — Sample Γ^{start}. (i) Guess Γ^{start}.

(ii) Simulate Γ^{start} forward using the existing $\mathcal{V}_\theta$.

(iii) Sample from the resulting trajectory of Γ^{start}, excluding burn-in.

Step 3 — Compute the outer Bellman residual at each sample.

For each Γ_i^{start} in the sample:

(i) Compute $(\mathcal{L}\mathcal{V}_\theta)\Gamma_i^{\text{start}}$ via lookahead:

[13] In machine learning, this operation is sometimes called "one-period lookahead."

(a) $V_i^{\text{end}} \leftarrow \mathcal{V}_\theta(\Gamma_i^{\text{start}})$.

(b) $\Gamma_i^{\text{end}} \leftarrow \Phi_\Gamma(\Gamma_i^{\text{start}}, V_i^{\text{end}})$.

(c) For a sample Z_k in the support of $\mathcal{D}$:

- $\Gamma_{i,Z}^{\text{start}\prime} \leftarrow \Omega(\Gamma_i^{\text{end}}, Z)$.
- $V_{i,Z}^{\text{end}\prime} \leftarrow \mathcal{V}_\theta(\Gamma_{i,Z}^{\text{start}\prime})$.
- $V_{i,Z}^{\text{start}\prime} \leftarrow \Phi_V(\Gamma_{i,Z}^{\text{start}\prime}, V_{i,Z}^{\text{end}\prime})$.

(d) Compute $\widehat{V^{\text{end}}}_i \leftarrow \mathbb{E}_Z[\, V_{i,Z}^{\text{start}\prime}\,]$ as a (possibly probability-weighted) mean of the $V_{i,Z}^{\text{start}\prime}$.

(ii) Compute the gradient of the outer Bellman residual, treating $\widehat{V^{\text{end}}}_i$ as a fixed label:[14]:

$$g_i \;\leftarrow\; \frac{\partial}{\partial\theta}\left\|\mathcal{V}_\theta(\Gamma_i^{\text{start}}) - \widehat{V^{\text{end}}}_i\right\|^2,$$

where $\|\cdot\|^2$ integrates the squared deviation across (a sample of) $x^{\text{end}} \in X^{\text{end}}$.

Step 4 — Update θ.

Take a gradient-descent step,

$$\theta \;\leftarrow\; \theta \;-\; \alpha \cdot \frac{1}{N}\sum_{i=1}^{N} g_i.$$

Step 5 — Evaluate convergence.

Repeat Steps 2–4 until

$$\left\|\mathcal{V}_\theta(\Gamma_i^{\text{start}}) - \widehat{V^{\text{end}}}_i\right\|^2 < \varepsilon$$

[14] That is, not computing gradients through the lookahead step