

Dynamic Programming in Economics on a Quantum Annealer*

Jesús Fernández-Villaverde

University of Pennsylvania, NBER, and CEPR

Isaiah Hull

BI Norwegian Business School and CogniFrame, Inc.

January 23, 2024

Abstract

We introduce a set of novel algorithms for solving dynamic programming problems in economics on a quantum annealer, a specialized device that performs combinatorial optimization. Quantum annealers attempt to solve an NP-hard problem by starting in a quantum superposition of *all* states and generating candidate *global* solutions in milliseconds, irrespective of problem size. In contrast to existing work on dynamic programming in quantum computing, our approach is 1) able to recover value and policy functions, 2) not limited by fundamental scaling bottlenecks, and 3) already implementable for small (but non-trivial) problems on existing quantum hardware. We illustrate how to apply our algorithms by solving the real business cycle model on existing quantum hardware.

Keywords: Computational Methods, Dynamic Equilibrium Economies, Quantum Computing, Quantum Annealing.

JEL codes: C63, C8, E37.

*We thank participants at numerous seminars and conferences for useful comments and suggestions. During part of the project, Isaiah Hull was affiliated with and had a financial interest in CogniFrame, Inc., which is a partner of D-Wave Systems, the producer of the quantum annealers used in the paper. The authors did not receive any support or compute time from D-Wave through this partnership. All access to D-Wave's quantum annealers was provided through a standard developer agreement.

1 Introduction

Quantum computing, i.e., the manipulation of quantum physical systems to undertake numerical computations, promises considerable speed-ups for many scientific problems, including those in economics. However, achieving a quantum speed-up is not simply a matter of translating a classical algorithm into its quantum equivalent.¹ This is not, in general, possible. Even when possible, it would not be expected to yield computational gains. Rather, achieving a quantum speed-up requires the construction of radically different algorithms that exploit quantum phenomena, such as superposition, entanglement, interference, and tunneling.

We tackle this challenge by proposing a set of novel algorithms to solve dynamic programming problems, such as those that appear in economics, on a quantum annealer (QA). This specialized quantum device performs combinatorial optimization using a physical process. A QA embeds the problem’s parameters in a quantum system that evolves to find its lowest energy configuration. This is equivalent to identifying the values of state variables that globally minimize a loss function (Farhi et al., 2000). QAs attempt to solve a problem that is NP-hard for a classical computer, start in a linear combination (quantum superposition) of all states, and return candidate solutions in milliseconds, irrespective of the problem size (Venegas-Andraca et al., 2018).

More concretely, our paper makes three key contributions:

1. The development of a set of novel algorithms for formulating dynamic programming problems as combinatorial optimization problems, which has broad applicability on both quantum and classical hardware.
2. The innovative use of state-of-the-art quantum annealing techniques, including reverse and inhomogeneous annealing. This enables us to execute our algorithm on a QA without hybridizing it into classical and quantum components, as is commonly done in the literature. Our approach can be applied to a broad class of iterative and multi-objective optimization problems.
3. The implementation and execution of the previous algorithm on existing quantum hardware rather than on a classical computer that simulates the behavior of a quantum computer.

These contributions are crucial because they address the limitations of QAs, which are not designed to solve dynamic programming problems. QAs do not naturally allow for iteration over time or across multiple objective functions and suffer from the quantum-to-classical bottleneck,

¹We use the term “classical” and “classically” to indicate that something is not quantum.

which severely limits how much classical information can be read out as the problem’s solution. We overcome these limitations with an approach that recovers policy and value functions, which are the objects of interest for most dynamic programming problems in macroeconomics, industrial organization, game theory, and labor economics. This contrasts with existing work in the quantum computing literature, which focuses primarily on recovering decision paths and other low-dimensional objects that are of interest in computer science problems.²

Our intention is to develop algorithms that are both 1) implementable for small problems on existing hardware and 2) free of scaling limitations that would prevent their eventual application to larger problems. Thus, we address two common problems in quantum algorithm development. The first is that theoretical work on quantum computing often demonstrates speed-ups using algorithms that may not be feasible to implement on actual quantum hardware in the foreseeable future. The second is that many proof-of-principle demonstrations on quantum hardware—which solve trivial problems on quantum hardware to demonstrate feasibility—have known scaling issues that will prevent them from ever being applied to challenging problems.

To accomplish our first goal, we apply our algorithms to solve the real business cycle (RBC) model on existing quantum hardware. Our approach is highly computationally efficient, achieving an execution time that is close to the theoretical minimum for the device used. Given the maturity of quantum technology as of the writing of this paper, it is not feasible to solve more complex models, such as a medium-scale New Keynesian model (even the RBC model may not have been solvable without our novel use of advanced annealing features).

To accomplish our second goal, we design algorithms that are free of known scaling problems. In particular, we use a parsimonious representation of the value and policy functions. Our approach ensures that our algorithms can be applied to problems of increasing size in the future. As quantum hardware matures in the coming years, the scaling properties of our algorithms will allow us to solve larger dynamic programming problems, potentially including some that are infeasible to solve classically.

To evaluate our results, we compare our quantum algorithms to i) the classical value function iteration (VFI) solution in [Aruoba and Fernández-Villaverde \(2015\)](#) (hereafter, [AFV](#)) and ii) parametric policy iteration (PPI). We pick VFI as the first comparison because it is straightforward to triangulate between VFI and any other solution algorithm in dynamic programming. Thus, readers may take their preferred algorithm (for instance, a parallelized projection method), compare it with VFI, and have a time comparison of this algorithm with our QA algorithm. We pick PPI as the second comparison because of its extreme speed in solving our RBC model. Even with the limitations of existing quantum technology, our quantum algorithms recover a solution that is sufficiently accurate for standard applications and has a competitive

²The traveling salesperson problem (TSP), for example, only requires the recovery of the shortest constrained path between two nodes in a graph.

execution time. Since the execution time does not depend on the size of the problem instance in the algorithms we construct, future hardware improvements will, in principle, permit us to solve problems of increasing size with *the exact same* execution time, something that classical computers will never be able to do. We interpret these results as demonstrating the *potential* of quantum hardware for economists, both in terms of designing specific algorithms and implementing them on QAs.

We differ from the existing literature on quantum dynamic programming in three key aspects. First, we use a QA rather than a universal quantum computer (UQC). As we discuss in Section 2, UQCs can be programmed to execute arbitrary quantum algorithms and allow for the proof of reductions in computational complexity. However, they are a less mature technology and, consequently, face more substantial barriers to the implementation of solution algorithms for non-trivial dynamic programming problems. Prior work on dynamic programming on UQCs has focused, instead, on the theoretical demonstration of quantum speed-ups and on proof-of-principle demonstrations for trivial problems.³ In contrast, our focus on QAs, coupled with the novel algorithms we construct, allows us to solve a standard RBC model on the current vintage of quantum hardware, rather than simulating not-yet-developed quantum hardware on a classical computer.

Second, we focus on problems of interest to economists and explore what is achievable with existing machines. Work on dynamic programming using quantum devices has centered almost exclusively on problems in physics and computer science that have naturally parsimonious solutions. In contrast, economists are typically concerned with recovering policy or value functions that live in high-dimensional hypercubes. This is a non-trivial difference since the quantum-to-classical bottleneck severely limits how much data can be read out of a quantum device. For instance, a QA with N qubits will start in an exponentially large quantum superposition state of dimension 2^N , but will eventually collapse into a classical state that contains only N classical bits of information. Thus, finding a way to encode the solution in N bits is a fundamental challenge for dynamic programming problems in economics, while it is usually not a primary concern in physics and computer science. In part, we overcome this encoding problem by using the parametric dynamic programming method introduced in [Benítez-Silva et al. \(2000\)](#) to reduce the space complexity of the solution.

Third, we show how to construct iterative algorithms that can be implemented on QAs, which no other work has accomplished to the best of our knowledge. A fundamental challenge of dynamic programming is that solution methods often require iteration, sometimes over alternating objective functions. QAs are not programmable in the way that classical computers or

³Even on UQCs, little theoretical progress had been made before [Ambainis et al. \(2019\)](#), who used the unstructured quantum search algorithm in [Grover \(1996\)](#), coupled with the computation of a partial dynamic programming table, to achieve a quadratic speed-up for problems that involve the selection of a subset of elements.

UQCs are programmable. That is, they can only embed a problem instance and then perform annealing with a limited set of parameters. Our approach addresses the problem of how to implement iterations by making use of two state-of-the-art features of QAs: reverse and inhomogeneous annealing. These two features widen the applicability of our work to any iterative algorithm, whether or not it is related to dynamic programming.

The rest of the paper is structured as follows. Section 2 gives an overview of quantum annealing. Section 3 introduces the benchmark RBC model we use as a testbed. In Section 4, we explain how to translate the model into a form that is solvable on a QA. Section 5 introduces hybrid and quantum algorithms for solving the model and compares the results to the benchmarks in AFV. Section 6 concludes with a discussion of the prospects of quantum annealing in economics. An Online Appendix adds further details.

2 Quantum Annealing

There are two primary models of quantum computing: universal quantum computing and quantum annealing. Universal quantum computers (UQCs) employ the “gate-and-circuit” model, in which quantum operations called gates are applied to quantum bits (qubits) to enable arbitrary computations. These gates must adhere to the principles of quantum physics, since they are performing an operation on a quantum system. Quantum circuits consist of gate sequences that implement a subroutine within a quantum system.

UQCs offer theoretical reductions in time and space complexity for various algorithms used in computational economics and econometrics (Hull et al., 2020). For instance, Ambainis et al. (2019) and Glos et al. (2021) prove theoretical quantum speed-ups for several classes of dynamic programming problem, including the traveling salesperson and vertex ordering problems. Experimental evidence also demonstrates that UQCs are capable of achieving a speed-up over classical computers for certain problems. Unfortunately, UQCs are technically demanding to construct, and their development has lagged behind specialized quantum devices, such as QAs.

In contrast, QAs do not offer theoretically provable reductions in computational complexity and cannot be programmed to execute arbitrary algorithms. Instead, QAs can only perform combinatorial optimization. This seemingly severe limitation has, however, enabled considerably faster experimental progress. At present, QAs have roughly 50 times as many qubits as UQCs.⁴

In the remainder of this section, we introduce the concept of quantum annealing, discuss the problems it can solve, explain its limitations, and identify problems in economics where it

⁴The largest UQCs are currently produced by IBM. Its *Eagle r1* quantum processor has 127 qubits. IBM has also recently introduced an exploratory *Osprey r1* processor with 433 qubits. In contrast, *Advantage*, the QA produced by D-Wave Systems that we use for our computations in this paper, has 5616 qubits.

might be applied. We start with a brief review of combinatorial optimization.

2.1 Combinatorial Optimization

Quantum annealing is a heuristic method for solving combinatorial optimization problems that may provide a substantial computational advantage over classical solution methods for specific cases (Zintchenko et al., 2015). For our discussion, we adopt a slightly modified version of the definition of a combinatorial optimization problem given in Venegas-Andraca et al. (2018).

Definition 1 (Combinatorial Optimization Problem). *Let E be a finite set with cardinality $|E| = n$, P_E the power set of E (hence $|P_E| = 2^n$), and C a loss function where $C : P_E \rightarrow \mathbb{R}$. The general setup of a combinatorial optimization problem is to find an element $\mathcal{P} \in P_E$ such that $C(\mathcal{P}) = \min_{\mathcal{P}_i \in P_E} \{C(\mathcal{P}_i)\}$.*

In game theory and mechanism design, problems involving the allocation of discrete objects often have a natural combinatorial form that fits the definition above and is conducive to using a QA. Problems that involve financial networks also align with this definition and can be solved on a QA, as shown in proof-of-principle exercises by Orús et al. (2019a,b).

Our focus will be on problems in economics that do not naturally fit this definition but are commonly reformulated as combinatorial problems to facilitate grid search or value function iteration. As the first example, consider a firm that minimizes a cost function, C , by choosing inputs, L and K , given technology $y^* = Y(L, K)$. If the problem does not permit an analytical solution, a common strategy is to search for the minimum of C over the Cartesian product of discrete grids, where $K \in \{K_0, K_1, \dots, K_{\bar{K}-1}\}$ and $L \in \{L_0, L_1, \dots, L_{\bar{L}-1}\}$. A solution to the problem is a pair of functions, $K^*(w, r, y)$ and $L^*(w, r, y)$, which yield optimal K and L choices given factor prices (w, r) and the target level of output y .

As a second example, consider the problem of an infinitely lived household that maximizes utility from consumption, subject to a budget constraint. This problem is often written as a dynamic program and solved using value function iteration. The solution is a discrete look-up table approximation of the optimal value function or policy rule for capital and consumption.

While these two examples can, in principle, be solved on a QA once we have reformulated them as combinatorial problems, the limitations of quantum computers will make this reformulation challenging. The primary difficulty, which we discuss in Subsection 2.5, is the quantum-to-classical bottleneck, which limits the amount of information that can be read out of a quantum computer, including a QA. The computational complexity of such problems also differs from that of more standard combinatorial optimization problems, which we discuss in Subsection 2.4.

2.2 Adiabatic Quantum Computing

Quantum annealing can be described as a *heuristic* or *finite temperature* implementation of adiabatic quantum computing (Venegas-Andraca et al., 2018; Shin et al., 2014), and its value as an optimization algorithm is justified through the quantum adiabatic theorem (Messiah, 1958; Ambainis and Regev, 2004). To explain these ideas, let us introduce the idealized concept of adiabatic quantum computing, which we will treat as the best-case scenario for quantum annealing that is unlikely to be achieved in practice.

Ambainis and Regev (2004) provide an informal definition of the *quantum adiabatic theorem*. A Hamiltonian expresses the total level of energy in a quantum system, where the eigenvalues of the operator correspond to the energy levels of the associated eigenvectors:

... if we take a quantum system whose Hamiltonian slowly changes from \mathcal{H}_1 to \mathcal{H}_2 , then, under certain conditions on \mathcal{H}_1 and \mathcal{H}_2 , the ground (lowest energy) state of \mathcal{H}_1 gets transformed to the ground state of \mathcal{H}_2 .

Farhi et al. (2000) showed that the previous fact can be used to construct an alternative model of quantum computing that exploits adiabatic evolution to identify global minima.

More concretely, one builds a quantum state that interpolates between two Hamiltonians: 1) a trivial Hamiltonian, \mathcal{H}_0 ; and 2) the problem Hamiltonian, \mathcal{H}_p . The system is initialized as \mathcal{H}_0 , but transitions to \mathcal{H}_p over an amount of time T , following the adiabatic evolution:

$$\mathcal{H}(t) = \left(1 - \frac{t}{T}\right) \mathcal{H}_0 + \frac{t}{T} \mathcal{H}_p.$$

The initial Hamiltonian, \mathcal{H}_0 , is specified to be trivial, such that we can identify the lowest level of energy (the ground state) analytically. In comparison, \mathcal{H}_p encodes a minimization problem of interest, where the energy level in the system corresponds to the loss associated with a state. Thus, \mathcal{H}_0 and \mathcal{H}_p can be interpreted as loss functions with ground states that correspond to global minima.

Returning to the description of the quantum adiabatic theorem, evolving from \mathcal{H}_0 to \mathcal{H}_p sufficiently slowly will ensure that the quantum system remains in the global minimum (ground state) in all periods. Thus, in principle, the quantum adiabatic theorem tells us how to find the global minimum of a combinatorial optimization problem.

Adiabatic quantum computing also allows us to compute a “speed limit” for the transition between \mathcal{H}_0 and \mathcal{H}_p to ensure the quantum system remains in its ground state. Unfortunately, for many problem types of size N , the speed limit takes the form $T = \mathcal{O}[\exp(\gamma N^\nu)]$, where γ and ν are positive parameters (Bapst et al., 2013; Lucas, 2014).

Hence, adiabatic quantum computing may require exponential time to transition between \mathcal{H}_0 and \mathcal{H}_p while remaining in the ground state. Consequently, it is unlikely that QAs will

solve hard global optimization problems in polynomial time. However, γ and ν could be considerably smaller than their corresponding values for the classical problem, enabling us to solve exponentially hard problems with larger input sizes. We will revisit this idea in Subsection 2.4.

2.3 Problem Encoding

Quantum annealing involves the evolution of a trivial Hamiltonian, \mathcal{H}_0 , configured in its ground state, into one that encodes the problem of interest, \mathcal{H}_p . If the transition happens adiabatically (very slowly), \mathcal{H}_p will also be in its ground state. We can then measure the quantum system and recover the global minimum.

QAs do not require us to specify \mathcal{H}_0 , since any trivial Hamiltonian with a known ground state (global minimum) will work.⁵ However, one must specify \mathcal{H}_p , a non-trivial task even for simple problems. This process is referred to as “problem encoding,” since it encodes the information about a minimization problem in the state of a quantum system. Next, we will discuss how such encodings are constructed, starting with the types of problems we can feasibly encode.

2.3.1 Conversion to a binary quadratic model

Subsection 2.1 highlighted that an optimization problem must be combinatorial to be solvable on a QA. This is because QAs must encode the structure of a minimization problem as a binary quadratic model (BQM).⁶ This can be done using the Ising or the quadratic unconstrained binary optimization (QUBO) model. The Ising model allows us to introduce the main concepts transparently and has an intuitive interpretation. The QUBO will be more convenient and, thus, our preferred choice for the rest of the paper. The QUBO model also abstracts from any physical system and, thus, does not require knowledge of physics.

The Ising model. The Ising model describes a system of atomic spins. For our purposes, it will suffice to understand that each spin is a physical system that can be configured in one of two states, which we will represent as +1 and -1. We can think of a spin as corresponding to a terminal qubit state in a QA.

We denote a vector of N spins as $s = \{s_0, s_1, \dots, s_{N-1}\}$. Each s_i has a bias, h_i , and each pair of spins, (s_i, s_j) , has a coupling, $J_{i,j}$, where $h_i, J_{i,j} \in \mathbb{R}$. The spins and couplings can be used

⁵In practice, an equal superposition of all states is typically used. Additionally, “reverse” annealing allows us to specify an initial classical state, which can be used to refine solutions by performing a local search.

⁶The term *quadratic* throughout this paper refers to the number of interactions between binary variables, not to the more familiar concept of a quadratic approximation in perturbation theory.

to define an Ising model:

$$\mathcal{H}(s) = \sum_{i=0}^{N-1} h_i s_i + \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} J_{i,j} s_i s_j.$$

Recall that \mathcal{H} encodes a loss function as the total energy level in a physical system. Thus, a terminal s vector that implies a higher \mathcal{H} also implies a higher loss. The biases and couplings are features of our optimization problem, which we encode in the system. The annealing process identifies an optimal configuration of spins to minimize the energy level.

	s_0	s_1	$\mathcal{H}(s)$
1	-1	-1	-1.0
2	-1	1	0.0
3	1	-1	1.6
4	1	1	-0.6

Table 1: Spin configurations and energy levels.

For concreteness, consider a problem where $N = 2$, $h = \{0.5, -0.3\}$, and $J = \{-0.8\}$. Table 1 enumerates all possible terminal s vectors and computes the total energy (loss) in the system. Combination 1, where $s = \{-1, -1\}$, is a global minimum for this problem, since $\mathcal{H}(-1, -1)$ is lower than for any other combination of spins. While this is trivial to prove for the $N = 2$ case, the classical time complexity of enumerating all possible combinations for an arbitrary s of length N is $\mathcal{O}(2^N)$.




Figure 1: A graphical representation of a binary quadratic model for the $N = 4$ case and with couplings between the pairs: $\{(s_0, s_1), (s_0, s_3), (s_1, s_2), (s_1, s_3)\}$.

BQMs, including the Ising model, have a natural graphical representation. Consider the case for a graph of size $N = 4$ with non-zero couplings between the pairs (s_0, s_1) , (s_0, s_3) , (s_1, s_2) , and (s_1, s_3) . Abstracting from biases and coupling magnitudes, we can represent this model with a graph, G , specified by vertices and edges, (V, E) . Figure 1 visualizes the graph, where $V = \{s_0, s_1, s_2, s_3\}$ and $E = \{(s_0, s_1), (s_0, s_3), (s_1, s_2), (s_1, s_3)\}$.

As we discuss in Section 4, the BQM implied by our problem must be mapped to a graph that is embedded in a QA. While the embedding is typically constructed in a classical pre-processing step – and, thus, is not a part of the solution – the problem’s difficulty and the embedding’s quality depend on the graph our problem implies (Lucas, 2014). Section 4 explicitly considers whether a BQM’s graph is suitable for a given QA topology.

Earlier, we presented the classical Ising model to enhance intuition. We described the Hamiltonian as an operator that expresses the level of energy in a system as a function of its state and the parameters of the system, which can be configured to encode an optimization problem. In addition, there is a version of the Ising model that allows for quantum phenomena and expresses the level of energy in the system.

To see this, we start with the standard initial state of a QA given by the Hamiltonian:

$$\mathcal{H}_0 = - \sum_{i=0}^{N-1} \sigma_i^x,$$

where $\sigma_i^x = I \otimes \dots \otimes \sigma^x \otimes \dots \otimes I$, I is a 2x2 identity matrix, and σ^x is the Pauli X matrix:

$$\sigma^x = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$

where the i indicates that σ^x is located in the i^{th} position in the sequence of tensor products: For the $N = 2$ case, for example, \mathcal{H}_0 would be specified as:

$$\mathcal{H}_0 = -(\sigma_x \otimes I + I \otimes \sigma_x) = \begin{bmatrix} 0 & -1 & -1 & 0 \\ -1 & 0 & 0 & -1 \\ -1 & 0 & 0 & -1 \\ 0 & -1 & -1 & 0 \end{bmatrix}.$$

The ground state corresponds to the eigenvector associated with the minimum eigenvalue of \mathcal{H}_0 . In this case, the characteristic polynomial of the expression for \mathcal{H}_0 is $\lambda^4 - 4\lambda = 0$, yielding the eigenvalues $\lambda = \{-2, 2, 0, 0\}$ and eigenvectors v :

$$v = \left\{ \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -1 \\ -1 \\ 1 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \\ 1 \\ 1 \end{bmatrix} \right\}.$$

The first eigenvector, $v_0 = [1, 1, 1, 1]$, is associated with the minimum eigenvalue ($\lambda_0 = -2$) and, thus, corresponds to the ground state of the system and the global solution to the problem.

With a suitable normalization, this vector represents a uniform superposition over all possible states.⁷ We have two qubits and two states in this case, so there are four possible classical configurations. A QA is typically prepared in this position (but for N qubits), since it has a trivially computable ground state. Performing measurement immediately would collapse the superposition, yielding one of the four possible states with equal probability.

During the annealing process, the Hamiltonian can be expressed as:

$$\mathcal{H} = \left(1 - \frac{t}{T}\right) \sum_{i=0}^{N-1} h_i \sigma_i^x + \frac{t}{T} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} J_{i,j} \sigma_i^z \sigma_j^z, \quad (1)$$

where the coefficients on the initial and problem Hamiltonians correspond to the annealing schedule. Again, we use σ_z^i as a shorthand for $\sigma_z^i = I \otimes \dots \otimes \sigma_z \otimes \dots \otimes I$, where σ_z is the Pauli Z matrix:

$$\sigma^z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

The Hamiltonian (1) contains *transverse fields*, the σ_i^x terms, which pull qubits toward a superposition of the +1 and -1 states. These fields are in tension with the σ_z^i terms in the problem Hamiltonian, which pull qubits toward a classical computational basis state.⁸

The QUBO model. The QUBO model is still a BQM, but it uses 0 and 1 states rather than the +1 and -1 states of the Ising model. A QUBO model allows equality and inequality constraints and does not use a quadratic approximation.⁹

The standard form of a QUBO problem is:

$$\mathcal{H}(x) = \sum_i^N Q_{i,i} x_i + \sum_{j<i}^N Q_{i,j} x_i x_j.$$

This Hamiltonian formulation emphasizes the relationship between the Ising model and QUBO problems in the context of quantum annealing.

Notice that $x = \{x_0, x_1, \dots, x_{N-1}\}$ is a vector of binary variables that can take on values of 0 and 1. The problem structure is embedded in Q , an upper-triangular matrix. We may rewrite the problem in terms of Q as $\mathcal{H}(x) = x^T Q x$.

⁷A superposition state is a linear combination of basis states. In this case, the system is in a uniform or equal superposition over the basis states.

⁸The eigenvectors of σ^z are $[1, 0]$ and $[0, 1]$, which correspond to +1 and -1 in the classical Ising model or 0 and 1 in the computational basis. The eigenvectors of σ^x are $[1/\sqrt{2}, 1/\sqrt{2}]$ and $[1/\sqrt{2}, -1/\sqrt{2}]$, which are equal superpositions of the classical Ising states.

⁹The QUBO model allows for the use of constraints, but they must be incorporated into the objective function. As we will see, this is also not limited to including higher-order terms, which can be added into a QUBO after *quadratization*.

Model conversion. Converting the constrained optimization problem implied by an economic model into a BQM is not trivial. Hence, as discussed next, we will take the intermediate step of mapping the objective function to a pseudo-Boolean function (PBF). We will then explain how to decompose the PBF into a QUBO and expand the objective function to include constraints.

2.3.2 Pseudo-Boolean functions

NP-hard problems must typically be converted to a PBF before they can be reduced to a QUBO (Venegas-Andraca et al., 2018). This includes most optimization problems in economics and finance that could be solved on a quantum annealer. A PBF takes the form $f : \mathcal{B}^N \rightarrow \mathbb{R}$, where \mathcal{B} is the Boolean domain.

Boros and Hammer (2002) show that each PBF has a unique representation as a multi-linear polynomial:

$$f(x) = \sum_{\mathcal{M} \subseteq \{1, \dots, N\}} \alpha_{\mathcal{M}} \prod_{i \in \mathcal{M}} x_i,$$

which can be reduced to a QUBO in polynomial time and with a polynomial bound on its size.

Thus, if we can map an optimization problem to a PBF, we can also reduce it to a QUBO and potentially solve it on a QA. This is not obvious since most challenging problems in economics do not have a clear mapping to a QUBO but may be reformulated as a PBF.

Partial example of PBF mapping. For illustration, consider an optimization problem that entails selecting next-period capital, k' , given the current capital stock, k , and a productivity shock, z . If we want to solve a discrete approximation to the problem, one approach is to define a grid for z and a shared grid for k and k' , where each maps an index to a value for a given variable. We can then construct a loss function that takes the model variables' indices, (i, j, m) , and returns a scalar-valued loss.

If we wanted to reconstruct the loss function as a QUBO, we would encounter two problems. First, the loss function requires us to specify the indices of three variables, but QUBOs are restricted to interactions between two binary variables. And second, the loss function takes integer-valued indices as inputs, but a QUBO specifies states in terms of binary variables.

An inefficient candidate solution to the second problem is to map each node in the grids for k , z , and k' to binary variables. That is, if k is at node j in its grid, then $x_{k_j} = 1$ and $x_{k_l} = 0$ for all $j \neq l$. Thus, we can represent the variable indices (i, j, m) from the original problem as $x_{k_i} x_{z_j} x_{k'_m} = 1$ in a PBF. However, we cannot include such a term in a QUBO since it consists of the product of three variables. One solution to this problem is to quadratize the term: that is, rewrite it in terms of quadratic and lower-order terms. Next, we show how to do this.

2.3.3 Degree reduction

Once we have mapped the initial problem to a PBF, the next step is to perform degree reduction until all non-quadratic terms are eliminated. A good quadratization strategy should yield a QUBO with the same global minimum as the PBF and minimize the number of auxiliary variables introduced.

Dattani (2019) and Dattani and Chancellor (2019) evaluate more than 30 methods of quadratization. In Online Appendix A, we discuss four methods that introduce either no auxiliary variables or the minimum number of auxiliary variables. In the text below, we introduce a method that can be applied in all circumstances but with less appealing properties. All methods are selected for their suitability for use on D-Wave’s *Pegasus* QA topology (the one employed by our application).

Reduction by substitution. A general strategy for performing degree reduction is to replace products of two variables in higher-order terms with an auxiliary variable and a quadratic constraint. Consider, for instance, the PBF:

$$\mathcal{H}(x) = 2x_1x_2x_3 + 4x_2x_3x_4 - 5x_2x_3x_5.$$

Since the product x_2x_3 is common to all monomials, performing the substitution $x_a = x_2x_3$ reduces all cubic terms to a quadratic form. However, it also requires the introduction of a quadratic constraint, yielding the modified form of the Hamiltonian:

$$\mathcal{H}'(x) = 2x_1x_a + 4x_ax_4 - 5x_ax_5 + \underbrace{\gamma(x_2x_3 - 2(x_2 + x_3)x_a + 3x_a)}_{x_a=x_2x_3}.$$

The penalty, $\gamma > 0$, is a hyperparameter that must be determined by the researcher.

While this approach provides a general-purpose strategy for degree reduction, it comes at a high cost: one auxiliary variable and one penalty parameter for each degree reduced. The penalty parameter is especially problematic since QAs are sensitive to the scale of the largest model parameter. Thus, introducing a sufficiently large γ to impose the constraint will weaken the relative strengths of couplings related to the problem itself.

2.3.4 Adding constraints

The *unconstrained* part of the term quadratic unconstrained binary optimization refers to constraints *outside* of the QUBO problem’s objective, which are not permitted. It is, however, possible to introduce constraints into the QUBO problem. For example, consider the case where we have the objective function $f(x) = x_1x_2x_3 + 3x_1x_3 + 2x_2$ and the constraint

$x_1 + x_2 = 1$. We can express this problem as an objective function with a penalty term, as in $\mathcal{H}(x) = x_1x_2x_3 + 3x_1x_3 + 2x_2 + \gamma(1 - x_1 - x_2)^2$.

2.4 Computational Complexity

Problems solved on a QA are mapped to a physical quantum Ising problem, which is NP-hard to solve exactly on a classical computer (Barahona, 1982). Additionally, the associated *decision problem* is NP-complete, which implies that the problem can be mapped to any other NP problem in a polynomial number of steps.

Consequently, if a QA could solve the Ising problem exactly in polynomial time, it would also be able to solve any NP problem in polynomial time. However, the theoretical and experimental evidence suggests that this is unlikely to be the case. Rather, it appears that QAs can find high-quality approximation solutions fast but can only slowly provide further improvements in the quality of the solution (Zintchenko et al., 2015; King et al., 2015).

QAs are unlikely to solve the Ising problem in polynomial time because the transition between the initial and problem Hamiltonian must be adiabatic to ensure that the system remains in the ground state. Faster transitions can lead to a jump in the energy level from the ground state to the first excited state, which can prevent the system from terminating in a state consistent with the global minimum at the end of the annealing process.

Thus, while the inability to perform the transition adiabatically prevents us from guaranteeing that we remain in the ground state, it does not rule out the possibility that we end in either the ground state or another low-energy state. In cases where the spectral gap remains large throughout the transition process, a QA may transition sufficiently slowly to ensure that this happens.

More precisely, the inability to guarantee that an annealing problem can be reduced to polynomial complexity might not preclude a substantial quantum advantage. Consider the case where a problem with exponential time complexity is reducible to one with polynomial time complexity, but in both cases, the computational resource requirements are sufficiently high that only small problem instances can be solvable. As depicted in Figure 2, the exponential time solution could solve larger problems. The line with circular markers represents an algorithm with exponential complexity, and the line with square markers represents an algorithm with polynomial complexity. The dashed line indicates the maximum permissible complexity given available computational resources. Thus, even if a QA requires exponential time, it can considerably speed up feasible problem cases.




Figure 2: Comparison of the computational complexity of algorithms with polynomial and exponential time complexity for different input sizes.

2.5 Limitations

Quantum computers enable us to execute algorithms that require less time and space than the best available classical algorithms for certain problem classes. But they also place unfamiliar limitations on the computations that can be performed. We discuss the most important limitations in this subsection.

Quantum-to-classical bottleneck. QAs use quantum superposition, which enables qubits to be in a linear combination of 0 and 1 states rather than in the 0 or 1 state, as with classical bits. Similarly, two bits on a classical computer can be in either the 00, 01, 10, or 11 configurations. In contrast, a pair of qubits on a QA can be in a linear combination of all four states simultaneously. More generally, for the N -qubit case, quantum superposition enables the system to be in a linear combination of 2^N states.

Superposition allows us to start a QUBO problem in a uniform linear combination over all 2^N states rather than in an individual state.¹⁰ During the annealing process, the system transitions from a quantum superposition state to a pure or basis state, which solves our problem. The QA then performs measurement and reads out classical bits corresponding to that state.

The fact that we cannot observe a quantum state directly and must instead sample from it using measurement is referred to as the *quantum-to-classical bottleneck*. Even if we create a quantum superposition over 2^N states, we can only retrieve N classical bits of information from

¹⁰If there are 2^N states, then the degree to which we will be in each of those states is $\frac{1}{2^N}$. If, for instance, we perform measurement on the system to see which state we are in, the superposition will collapse, and a random state will be chosen from a uniform distribution over the $\frac{1}{2^N}$ states.

each annealing step. Current state-of-the-art QAs have 5000 qubits, which means that each annealing step can return at most 0.625kB of classical information. Consequently, quantum annealing applications typically center around problems with a parsimonious solution.

Classical-to-quantum overhead. The time to perform the quantum annealing step is at least $5\mu s$ (microseconds) but is independent of the problem size. In addition to this, programming the quantum processing unit (QPU) takes an additional $9ms$ (milliseconds), and reading out the classical output of a project takes an additional $120\mu s$ (Venegas-Andraca et al., 2018). Thus, the total run time is $T = T_p + R(T_a + T_r)$, where T_p is the programming time, T_a is the anneal time, T_r is the readout time, and R is the number of repetitions.

We perform the anneal and readout R times but only program the QPU once. Furthermore, the programming step takes 72 times as long as the anneal and readout are combined. As a result, problems that need partitioning and solving using a hybrid algorithm that employs both a classical CPU and a QPU suffer from a substantial classical-to-quantum transfer overhead.

Since quantum annealing is a probabilistic process, it may yield a different result for each repetition. For this reason, we will typically perform multiple anneals and readouts for each problem ($R > 1$), allowing us to identify the frequency and energy level of different solutions. This will have the effect of reducing the importance of the transfer overhead.

Noise. Unlike classical computers, quantum computers cannot yet perform error correction or store information in memory for extended periods because of 1) the impossibility of copying quantum states (Wootters and Zurek, 1982); and 2) the natural tendency of quantum systems to decohere in nanoseconds or microseconds. In this respect, QAs have an advantage over universal quantum computers since the annealing algorithm is less affected by qubit decoherence (Childs et al., 2001; Albash and Lidar, 2015).

QPU topology. Once we have reduced our initial problem to a QUBO, we must embed it in the QPU’s graph. So, the graph topology will place another set of restrictions on the types of problems we can solve with a QA. Figure 3 shows a cell and a subgraph from the *Pegasus* topology, which is used in D-Wave’s *Advantage* QAs.

The primary topological considerations include the total number of qubits, the total number of couplers, and the number of couplers per qubit. Having fewer qubits means we can represent fewer binary variables in our QUBO and read out less classical information. Similarly, a lower total number of couplers means we will be less able to use *quantum entanglement* as a computational resource, which is needed to specify the strength and direction of relationships between variables.

In addition to the total number of couplers, the number of couplers per qubit will also be





Figure 3: Intra-cell and inter-cell connections on subgraphs of the 16x16x3 *Pegasus* topology. The left panel shows the connections within an eight-qubit cell. The right panel shows the inter-cell connections for an example cell embedded in a 4x4x3 subgraph. The figure was generated using the approach in [Dattani et al. \(2019\)](#).

important, as it will determine the extent to which *chaining* is needed. Chaining occurs when qubits are not directly coupled but have a quadratic term in the QUBO. It entails creating an indirect link by forcing the qubits between them to have the same state as one of the qubits in the quadratic term. Long chain lengths reduce the number of qubits available to represent the problem and force us to tune a *chain strength* hyperparameter.

2.6 Ideal Problems

Given the features of a QA, what constitutes an ideal problem that could demonstrate quantum advantage? The literature suggests that the following three problem attributes are important: 1) a large state space and parsimonious solution; 2) long, thin barriers between minima in the energy landscape (loss function); and 3) opportunities for reducing the loss through co-tunneling.

Large state-space and parsimonious solution. A problem that is solved directly on a QPU and cannot be partitioned into subproblems should have a solution that requires no more than N bits to express, where N is the number of qubits. D-Wave’s *Advantage* line of QAs, for example, have more than 5000 qubits and can output 0.625kB of classical information per anneal. An ideal problem would exploit quantum superposition by traversing a large state-space but should ultimately have a solution expressible using a small amount of classical information. A challenging equilibrium-finding problem, for instance, would be a good candidate.

Barriers between minima in the loss function. QAs exhibit *quantum tunneling*, where a qubit passes through a barrier in the energy landscape (Boixo et al., 2014; Albash et al., 2015b,a), rather than attempting to climb over it, as a classical optimization algorithm would do. Denchev et al. (2016) suggest that problems with loss functions that have long, thin barriers surrounding local minima are ideal candidates for tunneling. Figure 4 depicts the loss function for a problem that may permit a tunneling-induced speed-up. Classical solvers will get stuck in local minima, whereas QAs may tunnel through the barriers to the global minimum.




Figure 4: *Quantum tunneling* is depicted by the dashed line passing through the energy landscape.

Loss reduction through co-tunneling. Denchev et al. (2016) also exploit *co-tunneling* (Boixo et al., 2016), where multiple entangled qubits jointly tunnel through a barrier to a lower energy level (loss). If the qubits were not entangled, it would not be possible for all of them to tunnel through the barrier instantaneously. Ideal uses of co-tunneling involve areas in the loss function where changing one variable at a time in a particular direction will increase the loss, but changing all simultaneously will decrease it.

3 The Model

To show the potential of QA in economics, we solve the version of the real business cycle (RBC) model used in AFV, a well-known testbed for computation economics. Also, given the limitations of existing QAs at the time of writing this paper, this model is probably the limit of what can be implemented right now.

A social planner chooses a sequence of consumption c_t and capital k_{t+1} to solve:

$$\max_{\{c_t, k_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{\infty} (1 - \beta) \beta^t \ln(c_t)$$

subject to the resource constraint $c_t + k_{t+1} = z_t k_t^\alpha + (1 - \delta)k_t$, where β is the discount factor, z_t is a productivity shock, and \mathbb{E}_0 is the conditional expectations operator.

To match the comparison exercise in [AFV](#) and get a closed-form solution to benchmark our QA solution, we set $\beta = 0.95$, $\alpha = 0.33$, and $\delta = 1$. Recall that, with full depreciation, the optimal decision rules for the social planner are given as $c_t = (1 - \alpha\beta)z_t k_t^\alpha$ and $k_{t+1} = \alpha\beta z_t k_t^\alpha$. Also as in [AFV](#), z_t is a Markov chain with support $z_t \in \{0.9792, 0.9896, 1.0000, 1.0106, 1.0212\}$ and transition matrix:

$$\Pi = \begin{pmatrix} 0.9727 & 0.0273 & 0 & 0 & 0 \\ 0.0041 & 0.9806 & 0.0153 & 0 & 0 \\ 0 & 0.0082 & 0.9837 & 0.0082 & 0 \\ 0 & 0 & 0.0153 & 0.9806 & 0.0041 \\ 0 & 0 & 0 & 0.0273 & 0.9727 \end{pmatrix},$$

which approximates an AR(1) process for log productivity.

[AFV](#) write the model in the previous section as a Bellman equation:

$$v(k, z) = \max_{k'} \{ \ln(zk^\alpha - k') + \beta \mathbb{E}[v(k', z')|z] \}$$

and discretize the capital stock into 17,820 uniformly spaced points over $[0.5\bar{k}, 1.5\bar{k}]$, where \bar{k} is capital at the steady state. Then, they solve the problem using value function iteration (VFI). In C++, [AFV](#) report a computation time of 0.73 seconds (on a 2023 vintage PC, the running time would be around 0.4 seconds).

We cannot implement VFI on current QAs. State-of-the-art machines currently output 0.625kB of classical information per anneal, while the look-up table from the above dynamic programming problem requires 89,100 floating-point numbers, at least 285 times as much classical information, even if we use half-precision.

We get around this problem by using parametric dynamic programming (PDP), which allows for a representation of the solution that can be encoded in a small number of bits.¹¹ Furthermore, even if there were no quantum-to-classical bottleneck, PDP would still be preferable,

¹¹We will evaluate the value function at all productivity nodes and at a number of capital nodes proportional to the square root of the number of nodes in [AFV](#). This is already more than the number needed to identify the value function parameters. The parameters used to represent the value function in the quantum algorithms are selected from a set that has between 2^{14} and 2^{20} elements, depending on the algorithm.

since our intention is to eventually solve large problem instances on a quantum annealer. Even if we could output an exponential amount of classical information for a large problem instance, it would take an exponential amount of time and could not be stored on a classical computer.

In applications, PDP tends to yield solutions faster than VFI and is easier to scale up in terms of dimensionality of the state space, but it is often less stable (see, e.g., [Taylor and Uhlig, 1990](#), [Benítez-Silva et al., 2000](#), and [Sweeting, 2013](#)) and does not have a provably lower time complexity. Nonetheless, to ensure full compatibility of results, we will solve our PDP problem first on a classical computer to establish a clear benchmark.

4 Methods

We now describe how to set up the PDP problem associated with our RBC model, translate it into a PBF, quadratize the PBF into a QUBO, and execute the QUBO on a QA.

4.1 Setting up the PDP

We use a parametric policy iteration (PPI) formulation that follows [Benítez-Silva et al. \(2000\)](#) where both the policy and value functions have parametric forms to ensure they are expressed parsimoniously. PPI also typically yields solutions faster than discrete policy iteration but may have worse convergence properties.

We start by assuming functional forms for the value function and the consumption and capital policy rules:

$$v(k, z) = x_2 + x_3 \ln(y). \tag{2}$$

$$k' = x_1 y \tag{3}$$

$$c = (1 - x_1)y. \tag{4}$$

Then, we take the Bellman equation $v(k, z) = \max_{k'} \{\ln(c) + \beta \mathbb{E}[v(k', z')|z]\}$, and use equations (3) and (4) to rewrite it as:

$$v(k, z) = \max_{x_1} \{\ln[(1 - x_1)y] + \beta \mathbb{E}[v(x_1 y, z')|z]\}. \tag{5}$$

The functional forms are selected to align with the dynamic programming problem itself and are not reflective of specific requirements on the inputs of QAs.

In more complex models, one would need to specify richer functional forms for the value and decision rules. In particular, neural networks can be used to approximate the value function

and policy rules (Fernández-Villaverde et al., 2023). For example, we can write:

$$v(k, z) = x_0 + \sum_{i=1}^n x_i \phi(k, z, \mathbf{x}_{i,j})$$

where $\phi(k, z, \mathbf{x}_{i,j})$ is a suitable activation function indexed by a vector of parameters $\mathbf{x}_{i,j}$ and the x_i 's are the weights of the network. While this approximation requires more parameters than equations (3) and (4), it is conceptually straightforward and a simple extension of our approach.

PPI alternates between policy improvement and policy valuation steps. The policy improvement step uses equation (2) and is given by:

$$x_1^* = \underset{x_1}{\operatorname{argmin}} \{ -\ln[(1 - x_1)y] - \beta \mathbb{E}[\bar{x}_2 + \bar{x}_3 \ln(y') | z] \}.$$

Notice that \bar{x}_2 and \bar{x}_3 are taken as fixed parameters in this step, as indicated by the bar.

Re-arranging equation (5) yields a residual function. The policy valuation step minimizes the sum of the squared residuals by choosing x_2 and x_3 :

$$x_2^*, x_3^* = \underset{x_2, x_3}{\operatorname{argmin}} \{ \ln[(1 - \bar{x}_1)y] + \beta \mathbb{E}[x_2 + x_3 \ln(y') | z] - x_2 - x_3 \ln(y) \}^2.$$

To simplify notation, we denote the objective function for the policy improvement step as $g_p(x_1; \bar{x}_2, \bar{x}_3)$ and the objective function for the policy valuation step as $g_v(x_2, x_3; \bar{x}_1)$. Given some initial parameters $\{x_1, x_2, x_3\}$, PPI iterates on $g_p(x_1; \bar{x}_2, \bar{x}_3)$ and $g_v(x_2, x_3; \bar{x}_1)$ until achieving convergence. Algorithm 1 summarizes the previous steps.

Algorithm 1: Parametric Policy Iteration

- 1 Select a parametric form for the policy and value functions.
 - 2 Initialize the parameters of the policy function, x_1 , and value function, $\{x_2, x_3\}$.
 - 3 Update the policy function parameter: $x_1^* = \underset{x_1}{\operatorname{argmin}} \{g_p(x_1; \bar{x}_2, \bar{x}_3)\}$.
 - 4 Update the value function parameters: $x_2^*, x_3^* = \underset{x_2, x_3}{\operatorname{argmin}} \{g_v(x_2, x_3; \bar{x}_1)\}$.
 - 5 Repeat steps 3 and 4 until the convergence criterion is satisfied.
-

Coding Algorithm 1 on a classical computer is straightforward. Figure 5 plots the average results from 10 executions of Algorithm 1 on a classical computer. The horizontal axis shows the total execution time in microseconds. Each tick corresponds to an iteration. The vertical axis shows the parameter errors in absolute percentage deviations from their true values (which, in our version of the RBC model, we can compute analytically). The terminal errors for x_1 , x_2 , and x_3 are 1.37%, 0.77%, and 3.74% respectively. Convergence is achieved after 0.58 seconds,

which is faster than the 0.73 seconds required for the C++ benchmark VFI solution, but slightly slower than the 0.40 seconds required for the same benchmark solution on 2023-era hardware.

Therefore, our execution time comparisons with the AFV benchmarks later in the paper are not biased in favor of the QA because of the use of PPI instead of VFI; however, we will show how a QA can deliver solutions one order of magnitude faster than a classical computer, regardless of whether the latter uses PPI or VFI.




Figure 5: Policy and value function parameter errors as a function of average execution time over 10 runs. Convergence is achieved after two iterations.

Programming a quantum annealer to execute Algorithm 1 is more challenging since we must first re-express it as a QUBO problem. This transformation requires the intermediate step of formulating it as a pseudo-Boolean optimization (PBO) problem. The next two subsections show how we can implement these steps.

4.2 Translating the PDP into a PBO

Formulating the PDP problem as a pseudo-Boolean optimization (PBO) problem entails first mapping the parameters $\{x_1, x_2, x_3\}$ to binary variables. The simplest reformulation would discretize the parameter values, representing each node as a binary variable, and impose the condition that the sum of the binary variables equals one. However, this would be inefficient and soon exhaust the QA’s computational resources.

An alternative strategy, which we will use to represent $\{x_2, x_3\}$, re-expresses each variable using a standard binary encoding. This strategy allows us to represent N nodes x_ν in the state space using $\lceil \log_2(N) \rceil$ qubits:

$$x_\nu = s_\nu \sum_{j_\nu=0}^{J_\nu} 2^{j_\nu} x_{\nu, j_\nu}.$$

Notice that s_ν , where $\nu \in \{1, 2, 3\}$, is the scaling factor for variable x_ν , defined such that $x_\nu \in [0, s_\nu(2^{J_\nu+1} - 1)]$.

As a minimal example, consider the case for $\nu = 2$, where $s_2 = 1$ and $J_2 = 2$. If $x_{2,0} = 1$, $x_{2,1} = 0$, and $x_{2,2} = 1$, then $x_2 = 2^0 * 1 + 2^1 * 0 + 2^2 * 1 = 5$. If, instead, we want to set the maximum value to m , then we can impose the following scaling factor: $s_2 = \frac{m}{2^{J_2+1}-1}$. Raising J_2 increases the precision of x_2 for a fixed s_2 . Using just 19 qubits, we can construct a grid over x_2 values with more than 10^6 nodes that does not require any quadratic terms or constraints.¹²

A second, more substantial, challenge is representing the terms $\ln(x_1)$ and $\ln(1 - x_1)$ from Algorithm 1 on a QA. A naive approach might treat $\ln(x_1)$ and $\ln(1 - x_1)$ as separate variables and link them with a constraint. However, this approach is inefficient, as it requires the use of two sets of binary variables to represent $\ln(x_1)$ and $\ln(1 - x_1)$. It also requires us to invert $\ln(x_1)$ and $\ln(1 - x_1)$ and impose a quadratic constraint, which substantially increases the connectivity requirements and network size.

Instead, we use a single set of binary variables to represent x_1 and two sets of coefficients to transform it into $\ln(x_1)$ and $\ln(1 - x_1)$ using polynomial approximations:¹³

$$\begin{aligned} \ln(x_1) &\approx \left[a_0 + \sum_{j_1=0}^{J_1} a_1 2^{j_1} x_{1,j_1} + a_2 \sum_{j_1=0}^{J_1} \sum_{i_1=0}^{J_1} 2^{j_1+i_1} x_{1,j_1} x_{1,i_1} \right] \\ \ln(1 - x_1) &\approx \left[\tilde{a}_0 + \tilde{a}_1 \sum_{j_1=0}^{J_1} 2^{j_1} x_{1,j_1} \right], \end{aligned} \quad (6)$$

where $\{a_0, a_1, a_2\}$ and $\{\tilde{a}_0, \tilde{a}_1\}$ are the sets of coefficients.¹⁴

Then, we can exploit symmetry and the properties of binary variables (i.e., $x^2 = x$) to rewrite equation (6) as:

$$\ln(x_1) \approx \left[a_0 + \sum_{j_1=0}^{J_1} (a_1 2^{j_1} + a_2 2^{2j_1}) x_{1,j_1} + 2a_2 \sum_{j_1=0}^{J_1} \sum_{i_1 < j_1} 2^{j_1+i_1} x_{1,j_1} x_{1,i_1} \right].$$

This formulation is useful because it reduces the number of auxiliary variables that must be introduced to quadratize cubic and higher-order terms.

The use of quadratic terms does not require the introduction of additional binary variables. Hence, x_1 , x_2 , and x_3 can be represented using the same number of binary variables. Quadratic terms, however, require connectivity between qubits, a limited computational resource in QAs.

To fully characterize the PBO, we must re-express each of the terms in g_p and g_v as products and sums of binary variables. However, it is convenient to simplify first their expressions. We

¹²The difference in size between the coefficient on the smallest and largest terms will become problematic when using a binary variable encoding on a QA. Thus, it is still worthwhile to consider how much precision is needed for a given problem.

¹³Note that we are not performing a Taylor expansion around a point. Rather, we are representing $\ln(x_1)$ and $\ln(1 - x_1)$ in a binary form for permissible values of x_1 to satisfy the requirements of a QUBO.

¹⁴We have $[a_0, a_1, a_2] = [-0.10905, 0.57570, -1.38445]$ and $[\tilde{a}_0, \tilde{a}_1] = [-0.22278, -0.28375]$.

start by substituting on the definition of y' and collecting monomial terms:

$$g_p(x_1; \bar{x}_2, \bar{x}_3) = \{-\ln(y) - \ln(1 - x_1) - \beta\bar{x}_2 - \beta(\mathbb{E}[\ln(z')|z] + \alpha \ln(y))\bar{x}_3 - \alpha\beta \ln(x_1)\bar{x}_3\}.$$

Since y , \bar{x}_2 , and \bar{x}_3 are constants in g_p , we can minimize the simplified version of the objective function:

$$g_{p'}(x_1; \bar{x}_2, \bar{x}_3) = \{-\ln(1 - x_1) - \alpha\beta\bar{x}_3 \ln(x_1)\}.$$

This reformulation will have additional benefits. Most importantly, it eliminates the x_2 term. We will see below why this elimination makes it easier to construct algorithms executed entirely on a QPU.

We can now write down a PBO for the policy improvement step. Since $a_0 < 0$, $\tilde{a}_0 < 0$, and $\bar{x}_3 > 0$, we have $-(\tilde{a}_0 + \alpha\beta a_0 \bar{x}_3) > 0$. Consequently, removing the constant term might allow for negative energy states (loss function evaluations), which are problematic for one of the solution algorithms we will introduce in Section 5. Thus, we retain the constant and write:

$$x_{1,0}^*, \dots, x_{1,J_1}^* = \operatorname{argmin}_{x_{1,0}, \dots, x_{1,J_1}} \left\{ -(\tilde{a}_0 + a_0\alpha\beta\bar{x}_3) - \sum_{j_1=0}^{J_1} \left[(\tilde{a}_1 + a_1\alpha\beta\bar{x}_3)2^{j_1} + a_2\alpha\beta\bar{x}_3 2^{2j_1} \right] x_{1,j_1} - 2a_2\alpha\beta\bar{x}_3 \sum_{j_1=0}^{J_1} \sum_{i_1 < j_1} 2^{j_1+i_1} x_{1,j_1} x_{1,i_1} \right\}. \quad (7)$$

We next convert the policy valuation step into a PBO, starting with the expression:

$$g_v(x_2, x_3; \bar{x}_1) = \{\ln(y) + \ln(1 - \bar{x}_1) - (1 - \beta)x_2 + \zeta x_3 + \alpha\beta \ln(\bar{x}_1)x_3\}^2,$$

where we have defined $\zeta = \beta\mathbb{E}[\ln(z')|z] + (\alpha\beta - 1) \ln(y)$.

We now expand g_v :

$$\begin{aligned}
g_v(x_2, x_3; \bar{x}_1) = & \left\{ \gamma_0 + \gamma_1(\bar{x}_1) - \underbrace{2(1-\beta)(\ln(y) + \ln(1-\bar{x}_1))}_{\gamma_2(\bar{x}_1)} x_2 \right. \\
& + 2 \underbrace{\left[\ln(y)(\zeta + \alpha\beta \ln(\bar{x}_1)) + \ln(1-\bar{x}_1)(\zeta + \alpha\beta \ln(\bar{x}_1)) \right]}_{\gamma_3(\bar{x}_1)} x_3 \\
& - 2(1-\beta) \underbrace{\left[\zeta + \alpha\beta \ln(\bar{x}_1) \right]}_{\gamma_{23}(\bar{x}_1)} x_2 x_3 + \underbrace{(1-\beta)^2}_{\gamma_{22}} x_2^2 \\
& \left. + \underbrace{\left[\zeta^2 + 2\zeta\alpha\beta \ln(\bar{x}_1) + \alpha^2\beta^2 \ln(\bar{x}_1)^2 \right]}_{\gamma_{33}(\bar{x}_1)} x_3^2 \right\}. \tag{8}
\end{aligned}$$

Notice that $\gamma_0 = \ln(y)^2 - 2(1-\beta)\ln(y)$ and $\gamma_1(\bar{x}_1) = 2\ln(y)\ln(1-\bar{x}_1) + \ln(1-\bar{x}_1)^2$. We enclose \bar{x}_1 in parenthesis to indicate that the constant term γ_1 depends on \bar{x}_1 and will change after each policy improvement step.

Using the constants defined above, we rewrite equation (8) more compactly as:

$$g_v(x_2, x_3; \bar{x}_1) = \left\{ \gamma_0 + \gamma_1(\bar{x}_1) - \gamma_2(\bar{x}_1)x_2 + \gamma_3(\bar{x}_1)x_3 - \gamma_{23}(\bar{x}_1)x_2x_3 + \gamma_{22}x_2^2 + \gamma_{33}(\bar{x}_1)x_3^2 \right\}.$$

Finally, we define the PBO for the policy valuation step:

$$\begin{aligned}
x_{2,0}^*, \dots, x_{2,J_2}^*, x_{3,0}^*, \dots, x_{3,J_3}^* = & \underset{x_{2,0}, \dots, x_{2,J_2}, x_{3,0}, \dots, x_{3,J_3}}{\operatorname{argmin}} \left\{ \gamma_0 + \gamma_1(\bar{x}_1) \right. \\
& - \gamma_2(\bar{x}_1)s_2 \sum_{j_2=0}^{J_2} 2^{j_2} x_{2,j_2} + \gamma_3(\bar{x}_1)s_3 \sum_{j_3=0}^{J_3} 2^{j_3} x_{3,j_3} \\
& - \gamma_{23}(\bar{x}_1)s_2s_3 \sum_{j_2=0}^{J_2} \sum_{j_3=0}^{J_3} 2^{j_2+j_3} x_{2,j_2} x_{3,j_3} \\
& + \gamma_{22}s_2^2 \left(\sum_{j_2=0}^{J_2} 2^{2*j_2} x_{2,j_2} + \sum_{j_2=0}^{J_2} \sum_{i_2 < j_2} 2^{j_2+i_2+1} x_{2,j_2} x_{2,i_2} \right) \\
& \left. + \gamma_{33}(\bar{x}_1)s_3^2 \left(\sum_{j_3=0}^{J_3} 2^{2*j_3} x_{3,j_3} + \sum_{j_3=0}^{J_3} \sum_{i_3 < j_3} 2^{j_3+i_3+1} x_{3,j_3} x_{3,i_3} \right) \right\}. \tag{9}
\end{aligned}$$

4.3 Translating the PBF to a QUBO

Converting a PBF into a QUBO entails quadratizing all cubic and higher-order terms. This reduces products of three or more binary variables into expressions that contain products of no more than two binary variables without introducing approximation error.¹⁵ The policy improvement and valuation steps, as defined in equations (7) and (9), do not contain any such terms. Consequently, no quadratization is needed if x_1 enters g_v as a constant and x_2 and x_3 enter g_p as a constant, as they do in the hybrid algorithm we propose in Section 5.

If, however, we want to reduce the computational overhead by solving the problem on the QPU entirely, as we also do in Section 5, then we will need to introduce higher-order terms. Subsection 2.3.3 explained how to quadratize such terms. We will also give two examples here in the context of our problem.

We first consider the case where all terms in g_p are multiplied by the same binary variable, x_p . This will transform the quadratic terms into cubic terms in the last summation:

$$- \sum_{j_1=0}^{J_1} \sum_{i_1 < j_1} 2a_2\alpha\beta\bar{x}_3 2^{j_1+i_1} x_{1,j_1} x_{1,i_1} x_p.$$

Since $a_2\alpha\beta\bar{x}_3 2^{j_1+i_1} > 0$ for all \bar{x}_3 , j_1 , and i_1 , each term in the summation will be negative. Thus, we can use the NTR method described in Subsection 2.3.3 to quadratize the expression with a single auxiliary variable, x_{a,i_1,j_1} , for each term, yielding:

$$- \sum_{j_1=0}^{J_1} \sum_{i_1 < j_1} 2a_2\alpha\beta\bar{x}_3 2^{j_1+i_1} x_{a,i_1,j_1} (2 - x_{1,j_1} - x_{1,i_1} - x_p)$$

This step adds a total of $\frac{J_1(J_1-1)}{2}$ auxiliary variables to the policy valuation step.

Consider the case where x_1 does not enter g_v as a constant. This happens whenever we use an algorithm solved exclusively on the QPU and, thus, cannot specify x_1 as a parameter. The

¹⁵The quadratization methods we employ may introduce auxiliary variables, but do not affect the optimality of different state configurations and do not introduce approximation error.

term $\alpha^2\beta^2x_3x_1^2$, for example, becomes:

$$\begin{aligned} \alpha^2\beta^2x_3\ln(x_1)^2 &= \alpha^2\beta^2s_3s_1^2 \left\{ a_0^2 \sum_{j_3=0}^{J_3} 2^{j_3} x_{3,j_3} + a_0(1+a_1) \sum_{j_1=0}^{J_1} \sum_{i_3=0}^{J_3} 2^{j_1+i_3} x_{1,j_1} x_{3,j_3} \right. \\ &+ (2a_0a_2 + a_1^2) \sum_{j_1=0}^{J_1} \sum_{i_1=0}^{J_1} \sum_{i_3=0}^{J_3} 2^{j_1+i_1+j_3} x_{1,j_1} x_{1,i_1} x_{3,j_3} + \\ &+ 2a_1a_2 \sum_{j_1=0}^{J_1} \sum_{i_1=0}^{J_1} \sum_{l_1=0}^{J_1} \sum_{j_3=0}^{J_3} 2^{j_1+i_1+l_1+j_3} x_{1,j_1} x_{1,i_1} x_{1,l_1} x_{3,j_3} \\ &\left. + a_2^2 \sum_{j_1=0}^{J_1} \sum_{i_1=0}^{J_1} \sum_{l_1=0}^{J_1} \sum_{m_1=0}^{J_1} \sum_{j_3=0}^{J_3} 2^{j_1+i_1+l_1+m_1+j_3} x_{1,j_1} x_{1,i_1} x_{1,l_1} x_{1,m_1} x_{3,j_3} \right\}. \end{aligned}$$

The final three summations all consist of cubic or higher-order terms and, thus, require quadratization.¹⁶ Additionally, since $s_3 > 0$, all terms are positive, which precludes using NTR as a reduction method. We instead use the positive term reduction (PTR) method from [Boros and Gruber \(2014\)](#) and focus exclusively on the final summation in the expression, which consists of quintic terms. For each term, we will need three auxiliary variables:

$$\begin{aligned} \alpha^2\beta^2s_3s_1^2a_2^2 &\left[\sum_{j_1=0}^{J_1} \sum_{i_1=0}^{J_1} \sum_{l_1=0}^{J_1} \sum_{m_1=0}^{J_1} \sum_{j_3=0}^{J_3} x_{a_1,j_1,i_1,l_1,m_1,j_3} (3 + x_{1,j_1} - x_{1,i_1} - x_{1,l_1} \right. \\ &- x_{1,m_1} - x_{3,j_3}) + x_{a_2,j_1,i_1,l_1,m_1,j_3} (2 + x_{1,i_1} - x_{1,l_1} - x_{1,m_1} - x_{3,j_3}) + \\ &\left. x_{a_3,j_1,i_1,l_1,m_1,j_3} (1 + x_{1,l_1} - x_{1,m_1} - x_{3,j_3}) + x_{1,m_1}x_{3,j_3} \right]. \end{aligned}$$

We can also further reduce the number of auxiliary variables by exploiting symmetry and the properties of binary variables, as we have done previously.

5 Results

So far, we have explained how to express our PPI problem as separate QUBO problems for the policy and value function updates. But we are still not quite ready to solve our problem on a QA because Algorithm 1 alternates between policy and value function updates, an iterative structure QAs cannot handle.

To get around this problem, we proceed in four cumulative steps:

1. Subsection 5.1 runs the QUBO version of PPI, Algorithm 2, on a classical computer,

¹⁶Recall that a QUBO does not permit cubic or higher order terms. Quadratization can be used to reduce cubic and higher order terms to expressions of lower order terms without introducing approximation error.

iterating over value and policy function updates. This step allows us to compare the QUBO version of PPI with the original PPI in Algorithm 1, which we already ran on a classical computer (Figure 5).

2. Subsection 5.2 proposes a new quantum-classical hybrid algorithm, Algorithm 3, where we run one part of the PPI on a classical computer and the rest on a QA. Using these hybrid algorithms has become popular in the literature (Hull et al., 2020).
3. Subsection 5.3 considers a pure quantum algorithm, Algorithm 4, that exploits reverse and inhomogeneous annealing to perform a complete iteration of policy and value function updates within an anneal. It also eliminates the need to reprogram the QPU and iterates across anneals.
4. Subsection 5.4 modifies Algorithm 4 to perform multiple iterations within a single anneal and delivers Algorithm 5, which yields a candidate solution from each anneal. Algorithms 4 and 5 are the most innovative part of our paper from a technical perspective.

Let us go over each step now in order.

5.1 Classical Combinatorial PPI

First, we solve the PPI problem with a classical combinatorial algorithm that performs alternating iterations of policy improvement and valuation steps. This algorithm is the combinatorial version of Algorithm 1. The policy improvement step is solved analytically, whereas the policy valuation step is specified as a QUBO and solved exactly by searching over the discretized value function parameters. The details are given in Algorithm 2.

Algorithm 2: Classical Combinatorial Parametric Policy Iteration

- 1 Initialize the value function parameters, x_2 and x_3 .
 - 2 Compute x_1 analytically.
 - 3 Given the QUBO and x_1 , perform the policy valuation step by identifying the values of $\{x_{2,0}, \dots, x_{2,J_2}\}$ and $\{x_{3,0}, \dots, x_{3,J_3}\}$ that minimize the loss function.
 - 4 Map the solution to value function parameters: $\{x_{2,0}, \dots, x_{2,J_2}\} \rightarrow x_2$ and $\{x_{3,0}, \dots, x_{3,J_3}\} \rightarrow x_3$.
 - 5 Repeat steps 2-4 until the convergence criterion is satisfied.
-

Using an exact solution for x_1 provides us with a benchmark that has three advantages. First, it eliminates errors from the policy valuation step, yielding an approximate measure of errors arising from discretization. Second, it provides an upper bound for acceptable candidate solution run times since it is always better to use an exact solution when the execution time is

lower than the alternatives. And third, it can be used to evaluate how many iterations will be needed to achieve convergence in the hybrid and quantum algorithms.

We parameterize the algorithm by setting $J_2 = J_3 = 9$, which allows x_2 and x_3 to each take on 2^{10} discrete values, and requires a search over 2^{20} states on each policy valuation step. We select scaling factors, $s_2 = -0.035$ and $s_3 = 0.003$, to bound the value function parameters, $x_2 \in [0, 2x_2^*]$ and $x_3 \in [0, 2x_3^*]$, where x_2^* and x_3^* correspond to the true parameter values.




Figure 6: Policy and value function parameter errors as a function of average execution time over 10 runs. Convergence is achieved after two iterations.

Figure 6 plots the average results from 10 executions of the algorithm. The horizontal axis shows the total execution time in microseconds (only the run time varies across executions; the candidate solutions are identical). Each tick corresponds to an iteration. The vertical axis shows the parameter errors in absolute percentage deviations from their true values. The terminal errors for x_1 , x_2 , and x_3 are 0.04%, 1.24%, and 0.14%, respectively. Other parameter configurations yield lower errors for individual parameters but at the expense of higher errors for other parameters and a higher loss. Convergence is achieved after 3.33 seconds (two iterations), which is slower than the VFI solution in [AFV](#) and [Algorithm 1](#).

All run time comparisons here and below exclude initial pre-processing and terminal post-processing times. Since the benchmark we select is competitive and can be solved classically in under one second (see the results in [AFV](#)), we restrict the comparisons to tasks where there are substantive differences across algorithms. Removing fixed and polynomial time computational overhead common to all algorithms simplifies this comparison.

The classical computational steps are performed on a Linux machine with an 8-core Intel Xeon 2.00GHz processor and 50GBs RAM. The classical algorithms and classical components of hybrid algorithms were programmed in Python. The quantum algorithms and quantum components of classical algorithms were executed on D-Wave’s *Advantage* System 6.1, which has 5616 qubits, arranged in a lattice of 16 x 16 tiles known as a *Pegasus* (P16) graph. We

have released code that can be used to reproduce the results in the paper.¹⁷

5.2 Hybrid Quantum-Classical Algorithm

We next approach the PPI problem with a hybrid algorithm that performs alternating iterations of policy improvement and valuation steps. The policy valuation step is run on the QA, while the rest of the algorithm is run classically. The details are given in Algorithm 3. Notice that [C] indicates a step is performed classically and [Q] on a QA.

Algorithm 3: Hybrid Parametric Policy Iteration

- 1 [C] Initialize the value function parameters, x_2 and x_3 .
 - 2 [C] Compute x_1 analytically.
 - 3 [Q] Given the QUBO and x_1 , perform the policy valuation step, yielding $\{x_{2,0}, \dots, x_{2,J_2}\}$ and $\{x_{3,0}, \dots, x_{3,J_3}\}$ for each anneal.
 - 4 [C] Map the output of each anneal to value function parameters: $\{x_{2,0}, \dots, x_{2,J_2}\} \rightarrow x_2$ and $\{x_{3,0}, \dots, x_{3,J_3}\} \rightarrow x_3$.
 - 5 [C] Compute the mean values of x_2 and x_3 for the lowest energy level anneals.
 - 6 Repeat steps 2-5 until the convergence criterion is satisfied.
-

As before, we set $J_2 = 9$, $J_3 = 9$, $s_2 = -0.035$, and $s_3 = 0.003$, and use 100 anneals, each with a duration of $20\mu s$ (microseconds). We then select the 10% of anneals with the lowest energy levels and terminate the process after two iterations. The average energy level (loss) of the system tends to decline substantially over the first two iterations; however, subsequent steps have considerably smaller impacts on the loss. This aligns well with the classical results.

Notice that the hybrid algorithm contains quantum components that are affected by *true* randomness that arises in quantum systems. Consequently, the candidate solutions that the algorithm produces cannot be reproduced from an initial state. For this reason, we focus on the distribution of candidate solutions produced by 50 executions of the algorithm.

Table 2 reports the results. The error columns for x_1 , x_2 , and x_3 provide the absolute percentage deviation of the initial parameter values from their true values. The subsequent three columns show the total amount of computation time on the quantum processor (QPU Total), the amount of time needed to program the quantum processor (QPU Prog.), and the total execution time for both the classical and the quantum components (Total Time).

The hybrid algorithm tends to converge to a neighborhood of $\{x_1^*, x_2^*, x_3^*\}$ in two iterations. On average, it takes 1.76 seconds, 1.9 times faster than the classical combinatorial solution but 2.4 times slower than the fastest C++ implementation for VFI in AFV, which converged in 0.73 seconds. Importantly, only 3% of the hybrid algorithm’s execution time comes from QPU

¹⁷See the GitHub repository at <https://github.com/ijh85/quantum-dynamic-programming>, which contains code that can be used to verify the results in the paper.

	x_1 Error %	x_2 Error %	x_3 Error %	QPU Total	QPU Prog.	Total Time
Mean	5.70	1.34	9.63	5.23E+04	3.19E+04	1.76E+06
25th Percentile	3.12	0.58	5.33	5.23E+04	3.19E+04	1.53E+06
75th Percentile	7.74	2.03	13.22	5.23E+04	3.19E+04	1.60E+06
Std. Dev.	3.99	0.90	6.00	0.0E+04	0.00E+04	0.78E+06

Table 2: The initial values are $x_1 = 0.5$, $x_2 = -0.5$, and $x_3 = 0.5$. All times are given in microseconds. Errors for the policy and value function parameters are given as absolute percentage deviations from their true values. Summary statistics are reported from 50 executions of the same hybrid program. The quantum part of each iteration uses 100 anneals and selects the 10% with the lowest energy levels.

computation. And only 39% of that is attributable to annealing. The rest is overhead from the QPU programming step, a computational cost that precedes a sequence of anneals for a given problem.

Thus, moving to a non-hybrid quantum algorithm could substantially improve computational time. An algorithm requiring the QPU to be programmed only once could reduce the total execution time from 1.76 seconds to 0.026 seconds.

Finally, notice that the hybrid algorithm yields less precise results on average than the classical benchmarks. The terminal value of x_3 , for instance, deviates from x_3^* by 9.63% on average. This is in line with the literature, which finds that QAs can produce high-quality approximations to the global solution fast but require a substantial amount of time to refine the solution further (Zintchenko et al., 2015; King et al., 2015). If higher precision is needed, the QA solution can be used to warm-start a classical algorithm.

5.3 Multi-Anneal Quantum Algorithm

Motivated by the previous results, we construct an algorithm that eliminates the classical components of the algorithm and requires the QPU to be programmed only once. More concretely, we specify a single QUBO, along with an annealing schedule, that can iteratively optimize the two components of the objective function. To so do, we introduce two novel computational strategies, both of which use recent developments: reverse annealing and inhomogeneous annealing.

Reverse annealing initializes the QPU in a classical state rather than an equal superposition over the entire state space. It then opens up a superposition state weighted toward the initial classical state. Specifying a full reversal is equivalent to performing a standard forward anneal, whereas specifying a small, partial reversal is likely to yield a solution close to the initial classical state. In standard usage, reverse annealing is applied to refine candidate solutions by searching within a specified neighborhood (Pelofske et al., 2020). The extent of the reversal determines the size of the neighborhood.

Inhomogeneous annealing supplies the QPU with separate offset values for each qubit. A positive offset indicates that the specified qubit should be annealed ahead of the global schedule, while a negative value indicates that it should be annealed afterward. Inhomogeneous annealing has been applied to improve alignment between multiple qubits in the annealing process (Adame and McMahan, 2020).

We use reverse and inhomogeneous annealing to facilitate iteration over two separate objective functions, which differs from standard practice. Algorithm 4 performs this iteration across anneals. Algorithm 5, which we will introduce in the next subsection, demonstrates how to iterate within an anneal. In both cases, the QPU is only programmed once.

Algorithm 4 describes the multi-anneal quantum routine. We parameterize the algorithm by setting $J_1 = J_2 = J_3 = 6$, which allows x_1 , x_2 , and x_3 to each take on 2^7 discrete values.

Algorithm 4: Quantum Parametric Policy Iteration

- 1 Quadratize the policy improvement and valuation PBFs separately, reconcile the resulting QUBOs, and then merge them.
 - 2 Initialize the classical states of policy function variables, $\{x_{1,0}, \dots, x_{1,J_1}\}$, and value function variables, $\{x_{2,0}, \dots, x_{2,J_2}, x_{3,0}, \dots, x_{3,J_3}\}$.
 - 3 Initialize variables that activate and deactivate loss components, $x_p = 0$ and $x_v = 0$.
 - 4 Set an inhomogeneous annealing schedule that first anneals $\{x_{1,0}, \dots, x_{1,J_1}\}$ and x_p , and then anneals $\{x_{2,0}, \dots, x_{2,J_2}, x_{3,0}, \dots, x_{3,J_3}\}$ and x_v .
 - 5 Set a reverse annealing schedule that applies to all qubits.
 - 6 Perform \mathcal{N} anneals without reinitializing the original classical state.
 - 7 Select the results from the anneal with the lowest energy level.
-

Step 1 uses a non-standard approach to quadratize the PBF into a QUBO. This entails quadratizing $x_p g_p$ and $x_v g_v$ separately, reconciling the two resulting QUBOs, and merging them. This approach ensures that qubits are otherwise inactive and annealed with the correct group. Steps 2 and 3 initialize the QA in a classical state where both components of the objective function are inactive (i.e., $x_p = 0$ and $x_v = 0$). In steps 4 and 5, the reverse anneal opens up a superposition for $\{x_{1,0}, \dots, x_{1,J_1}\}$ and x_p , allowing for optimization over the policy function and with the policy function parameters, but not over the value function or with the value function parameters. As the anneal progresses through its schedule, the qubit offsets specify that a superposition should be opened for $\{x_{2,0}, \dots, x_{2,J_2}, x_{3,0}, \dots, x_{3,J_3}\}$ and x_v , allowing for optimization over the value function and with value function parameters. The modified objective function is given by:

$$g = x_p g_p(x_1; \bar{x}_2, \bar{x}_3) + x_v g_v(x_2, x_3; \bar{x}_1)$$

Two additional details complete the algorithm. First, the QA should not be set to reinitialize after each anneal. This will ensure that each anneal in a sequence starts in the terminal classical state of the previous anneal, allowing for the repetition of both the policy improvement and the

valuation steps. Second, the problem and annealing schedule should reliably yield $x_p = 0$ and $x_v = 0$ as the terminal classical states. This is because x_p must converge to 0 to ensure that the policy function component of the objective is inactive (or small) during the policy valuation step. It also ensures that the system starts the next anneal with $x_p = 0$ and $x_v = 0$. This can always be achieved by adding a positive bias term to one of the components of the objective function. The term’s magnitude can be increased until anneals reliably yield $x_p = 0$ and $x_v = 0$.





Figure 7: Stylized depiction of the annealing schedule.


Figure 7 provides a stylized depiction of the annealing schedule. The policy function parameters and the objective function component are annealed first (dashed line), followed by the value function and its parameters (continuous line). In both cases, the parameters start in the classical state and are then reverse annealed and forward annealed. The annealing activity in each step is depicted as a weighted graph in Figure 8. A gray node indicates that a qubit is in a classical state, whereas a black node indicates that it is a superposition state undergoing an anneal. Black edges indicate that at least one of the connected nodes is being annealed. Notice that the policy improvement step has more annealing activity, since the quadratic terms used to approximate $\ln(x_1)$ require connectivity.

Table 3 provides summary statistics for 50 executions of Algorithm 4, each performing 20 anneals. For each QPU execution, we select the parameter values associated with the lowest-energy anneals for each objective function.¹⁸ The total QPU computation time is now reduced to 0.0217 seconds, approximately 3% of the 0.73 second computation time of the VFI solution

¹⁸The terminal energy levels of the annealer are not a reliable metric for comparing anneals, since we use a non-standard objective function that is zeroed out over the anneal schedule. Instead, we reconstruct the two components of the objective function, as described in Subsection 5.3.



Policy Improvement Graph



Policy Valuation Graph

Figure 8: State of the problem graph during the policy improvement and policy valuation steps.

using C++ in [AFV](#) or 0.66% of the computation time of the classical combinatorial algorithm in Subsection 5.1.

	x_1 Error %	x_2 Error %	x_3 Error %	QPU Total	QPU Programming Time
Mean	12.13	5.82	14.50	2.17E+04	1.60E+04
25th Percentile	9.62	3.33	7.40	2.17E+04	1.60E+04
75th Percentile	14.82	7.21	20.22	2.17E+04	1.60E+04
Standard Deviation	3.58	3.62	9.26	0.00E+00	0.00E+00

Table 3: We use a reverse annealing parameter of 0.00, which corresponds to a 100% reversal. All times are given in microseconds. Summary statistics are reported from 50 QPU executions of the same program, each of which performs 50 anneals.

While the run time is reduced, the mean policy and value function parameter errors across the 50 QPU executions increase slightly relative to the hybrid algorithm. Also, iterating across anneals without resetting the QPU results in high correlation across candidate solutions within a QPU execution. Consequently, we execute the QPU multiple times to get independent solution candidates. In the next subsection, we propose a refinement that transforms each anneal into an independent candidate solution. This allows us to reliably obtain high-quality results from programming and executing the QPU once.

5.4 One-Shot Quantum Algorithm

Next, we propose a *one-shot* quantum PPI algorithm, Algorithm 5. In contrast to Algorithm 4, Algorithm 5 allows for multiple iterations over the policy and value functions within a single anneal. In addition, it produces terminal energy levels that can be directly compared across anneals, allowing us to reliably select the candidate solutions with the lowest associated losses. This is not possible with Algorithm 4 because x_p and x_v converge to zero over the annealing step. Relative to Algorithm 4, Algorithm 5 has two drawbacks: it requires longer annealing schedules and longer pauses between anneals to reinitialize the state.

Algorithm 5: One-Shot Quantum Parametric Policy Iteration


- 1 Quadratize the policy improvement and valuation PBFs separately, reconcile the resulting QUBOs, and then merge them.
 - 2 Initialize random classical states for the policy and value function variables:
 $\{x_{1,0}, \dots, x_{1,J_1}, x_{2,0}, \dots, x_{2,J_2}, x_{3,0}, \dots, x_{3,J_3}\}$.
 - 3 Initialize (de-)activation variables in classical zero states: $x_v = x_p = 0$.
 - 4 Set an inhomogeneous annealing schedule to activate the policy improvement (g_p) and policy valuation (g_v) functions in sequence, and to anneal their associated variables \mathcal{C} times.
 - 5 Set the global annealing schedule to perform a full reversal.
 - 6 Perform \mathcal{N} anneals and reinitialize the state after each.
 - 7 Compute alternative measures of terminal state energy for each anneal.
 - 8 Return average parameter values over the lowest energy anneals.
-

By allowing for iteration within the annealing step, we can perform a full reverse anneal and search the entire state space, rather than searching just a neighborhood around the terminal state of the previous anneal. Consequently, the terminal state of each anneal may be treated as a candidate solution rather than an iteration. Additionally, since no information needs to be retained across anneals, we may reinitialize the state to reduce the correlation across candidate solutions.


Algorithm 5 contains two novel components: 1) an inhomogeneous reverse annealing schedule that permits multiple iterations *within* an anneal; and 2) a post-processing routine that boosts the quality of solutions. Let us discuss each of them before executing the algorithm and discussing the results.

Annealing schedule. Figure 9 illustrates the annealing schedule we propose for the one-shot algorithm, which contains reversals and is inhomogeneous.


The top panel illustrates the case where qubits are partitioned into two groups annealed separately but within a single *cycle*. We will use the term cycle to refer to a subschedule where all qubits are annealed once. The top panel contains one such subschedule and, thus, one cycle,



One Cycle Experimental Annealing Schedule



Two Cycle Experimental Annealing Schedule



Three Cycle Experimental Annealing Schedule

Figure 9: Experimental annealing schedules for the case of 1, 2, and 3 cycles. In all cases, we adopt a reverse, inhomogeneous anneal. Group 1 contains qubits $\{z_0\}$ for \mathcal{H}_s and $\{z_0, z_2\}$ for \mathcal{H}_c . Group 2 contains qubits $\{z_1\}$ for \mathcal{H}_s and $\{z_1, z_3\}$ for \mathcal{H}_c . All times are given in microseconds (μs).

whereas the middle and bottom panels contain two and three cycles since they are repetitions of the subschedule in the top panel.

Our proposed algorithm deviates from the original conception of quantum annealing, which was structured around a simple forward anneal (Farhi et al., 2000). Furthermore, existing applications of reverse annealing typically focus on solution refinement (Pelofske et al., 2020),

whereas we exploit reverse annealing to construct iterative algorithms. To demonstrate that our Algorithm 5 performs as intended, we conduct two experiments in a simpler setting and report the results in Online Appendix B.

As demonstrated in the simulations, cycling over repetitions of the same annealing subschedule allows us to emulate an iterative algorithm within a single anneal. This suggests that the annealing schedule in Algorithm 5, coupled with the (de-)activation mechanism, could provide the foundation for a one-shot quantum algorithm.

Finally, it is worth explaining why the initialization matters, since we use a full anneal reversal. As shown in Figure 9, the qubits are partitioned into two groups, which are annealed separately in each cycle. As such, one group will be in a classical state while the other is reversed into a full superposition state prior to the forward annealing step. The initialization aims to pin down the classical states of Group 2 qubits during the first round of annealing for Group 1. However, Group 1 qubits are immediately reversed into a uniform superposition and, thus, may be initialized in any state.

Post-processing. Algorithm 5 yields a candidate solution from each anneal. However, there are often substantial differences in the quality of the candidate solution. This is captured by the system’s terminal energy level, which is used in more standard problems to identify the best candidate solution. As discussed in conjunction with Algorithm 4, these energy level readouts are not a reliable measure of solution quality for our problem since the weights of different components of the Hamiltonian are shrunk during parts of the annealing schedule.

Thus, we modify the post-processing routine to improve the interpretability of energy levels. We start by recomputing the energy level in each terminal state without applying the de-activation qubits. We refer to this measure as the *unadjusted loss* since it provides an accurate measure of the loss but makes no further transformations. As we will see, however, this measure of the loss is still imperfect since it combines losses from both objective functions and is dominated by errors in x_2 . In practice, the unadjusted loss will be helpful for evaluating the quality of a solution for x_2 , but not for x_1 and x_3 .

To account for these issues, we construct another measure of loss, which we refer to as the *minimum loss*. This measure isolates the impact of errors in parameter x_j on the component of the loss function associated with x_j . It does this by fixing the values of the other parameters at their true values. For example, for the policy function parameter, we would compute the minimum loss associated with a candidate solution, x_1^c , as $g_p(x_1^c; x_2^*, x_3^*)$. This accomplishes two things. First, it reduces the variance in energy levels across anneals attributable to errors in the other parameters. And second, it ensures that the correct value of the parameter of interest (e.g., x_1^*) will be consistent with the lowest possible energy level.

Unfortunately, it is not generally feasible to compute the minimum loss, since we would not

typically know the true values of other parameters. Thus, we will use instead the *adjusted loss*, which is feasible to construct and approximates the minimum loss. We compute this measure by dropping the x_2 terms from x_p , where they only enter additively. We then compute the loss separately for each parameter, fixing the values of all other parameters at their mean values over the lowest energy anneal solutions. This eliminates much of the variance across anneals in g_p that is attributable to errors in x_2 and x_3 .

	x_1 Error %	x_2 Error %	x_3 Error %
Unadjusted Loss	-0.01	0.92	0.05
Minimum Loss	0.97	0.96	0.97
Adjusted Loss	0.97	0.96	0.97

Table 4: Correlations between the percentage error in the solutions for the policy and value function parameters.

As a final step, we retain the subset of anneals with the lowest adjusted loss for a given parameter and compute the mean value of the parameter over those anneals. Table 4 compares the correlations among parameter errors and the values of the three losses above. As discussed previously, there is a strong relationship between the unadjusted loss and the error in x_2 , almost as strong as the minimum loss. In comparison, the correlations between the errors in x_1 and x_3 and the unadjusted loss is weak. This suggests that the unadjusted loss is not a helpful metric. The minimum loss, however, is strongly correlated with the parameter errors, suggesting that the elimination of errors in x_2 and x_3 , for example, would enable us to identify anneals that contain useful information about x_1^* . The adjusted loss appears to be largely sufficient for this task, yielding a correlation coefficient of 0.97 for x_1 , 0.96 for x_2 , and 0.97 for x_3 .

Figure 10 further illustrates the benefits of post-processing anneals using the adjusted loss rather than the unadjusted loss. Each subfigure shows a binned scatterplot of loss values against absolute parameter value errors. The adjusted and unadjusted losses are roughly equivalent for x_2 . However, for x_1 and x_3 , the low values of the adjusted loss are highly informative about true parameter values, while the low values of the unadjusted loss are much less so.

Results. Table 5 reports descriptive statistics for 50 QPU executions of Algorithm 5. In each case, we perform 200 anneals. Within each anneal, we execute three cycles of the annealing subschedule, as illustrated in the third panel of of Figure 5. The errors for the policy and value functions are lower on average than those generated by the multi-anneal quantum algorithm. Also, the errors for x_1 and x_3 are lower on average than those generated by the hybrid algorithm. The total execution time on the QPU is 0.065 seconds, which is 3.8 times as long as the multi-anneal quantum algorithm, but 3.7% of the run time of the hybrid algorithm and 8.9% of the run time of the C++ implementation of VFI in AFV, which was 0.73 seconds.




Figure 10: Parameter error in absolute percentage deviations with levels of the *unadjusted loss* and *adjusted loss*. This is visualized using a binned scatter plot and fitted line for the results from 200 anneals. The vertical axis uses a log scale because the absolute error typically increases sharply at low levels of the *adjusted loss*.

The run time increases relative to the multi-anneal algorithm for three reasons. First, the annealing schedule contains multiple cycles of reverse and forward annealing, which increases the time per anneal from $23\mu s$ to $115\mu s$. Second, the classical state needs to be reinitialized between anneals to increase the independence of candidate results, adding overhead to each anneal. And finally, we use 200 anneals to identify independent candidate solutions, rather than iterating over a few highly correlated anneals.

To provide further intuition for the size of the errors, we conduct a 10-period simulation exercise in which we compute the consumption path following a negative productivity shock.

	x_1 Error %	x_2 Error %	x_3 Error %	QPU Total	QPU Programming Time
Mean	2.98	1.71	4.52	6.52E+04	1.59E+04
25th Percentile	1.70	1.08	1.92	6.52E+04	1.59E+04
75th Percentile	2.49	2.01	3.98	6.52E+04	1.59E+04
Standard Deviation	3.48	0.87	5.47	0.00E+04	0.00E+04

Table 5: We use a reverse annealing parameter of 0.0, which corresponds to a 100% reversal. All times are given in microseconds. All anneals use three cycles. Summary statistics are reported from 50 QPU executions of the same program, each of which performs 200 anneals.

Figure 11 shows the paths implied by the analytical solution and the one-shot quantum algorithm solution. The difference is negligible in each period but is most pronounced immediately after the shock.



Figure 11: Consumption simulation paths generated from the true parameters and the one-shot solution parameters. In both cases, we simulate the 10-period consumption response to a negative productivity shock in the first period.

In summary: Algorithm 5 provides a quantum solution that fully explores the state space, reduces correlation across the terminal states, and generates higher-quality solutions than the multi-anneal algorithm. It also provides an order of magnitude speed-up over classical benchmarks but is slower than the multi-anneal algorithm.

6 Conclusion

We introduce a collection of algorithms for solving dynamic programming problems and other computational problems with an iterative structure on a quantum annealer (QA). Our approach is capable of recovering value and policy functions, and can be implemented on existing

quantum hardware for small (but non-trivial) dynamic programming problems in economics. By construction, our algorithms do not suffer from fundamental scaling limitations, which will enable their application to models of increasing size as the quantum hardware matures.

We have illustrated the feasibility of implementing our algorithms on existing hardware by solving the RBC model in [AFV](#). Through the construction of two fully quantum algorithms that do not require hybridization, we achieved run times that approach the theoretical minimum for our computational task on a QA. While we do observe reduced accuracy relative to classical (non-quantum) solutions, it is sufficiently high for standard applications of the RBC model.

We have also discussed other uses of quantum annealing for dynamic programming. Quantum annealing could, for instance, be used to warm-start high-dimensional dynamic programming problems by mapping the parametric solution produced by the QA to initial values for the dynamic programming table employed by a more familiar classical algorithm.

Quantum annealing can also identify the average or maximum error size for a dynamic programming problem on a large grid. This could be useful if we have a candidate solution but cannot evaluate its accuracy in regions far from the steady state without significant computational expense. It could also identify features of the policy or value function that lead to performance degradation if not accounted for explicitly.

In summary, we hope to have convinced the reader that quantum computing opens many new doors in solving economic models. QAs are already good enough to solve some problems of interest in economics and are competitive with existing VFI-based benchmarks in the literature and with PPI. Much research, though, remains to be done.

References

- Adame, J. I. and McMahon, P. L. (2020). Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance. *Quantum Science and Technology*, 5(035011).
- Albash, T. and Lidar, D. A. (2015). Decoherence in adiabatic quantum computation. *Physical Review A*, 91(6).
- Albash, T., Rønnow, T., Troyer, M., and Lidar, D. (2015a). Reexamining classical and quantum models for the D-wave one processor. *The European Physical Journal Special Topics*, 224(1):111–129.
- Albash, T., Vinci, W., Mishra, A., Warburton, P. A., and Lidar, D. A. (2015b). Consistency tests of classical and quantum models for a quantum annealer. *Phys. Rev. A*, 91:042314.
- Ambainis, A., Balodis, K., Iraids, J., Kokainis, M., Prūsis, K., and Vihrovs, J. (2019). *Quantum Speedups for Exponential-Time Dynamic Programming Algorithms*, pages 1783–1793. ACM-SIAM.
- Ambainis, A. and Regev, O. (2004). An elementary proof of the quantum adiabatic theorem. arXiv:quant-ph/0411152.
- Aruoba, S. B. and Fernández-Villaverde, J. (2015). A comparison of programming languages in macroeconomics. *Journal of Economic Dynamics and Control*, 58:265–273.
- Bapst, V., Foini, L., Krzakala, F., Semerjian, G., and Zamponi, F. (2013). The quantum adiabatic algorithm applied to random optimization problems: The quantum spin glass perspective. *Physics Reports*, 523(3):127–205.
- Barahona, F. (1982). On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241–3253.
- Benítez-Silva, H., Rust, J., Hitsch, G., Pauletto, G., and Hall, G. (2000). A comparison of discrete and parametric methods for continuous-state dynamic programming problems. *Computing in Economics and Finance 2000 24*, Society for Computational Economics.
- Boixo, S., Rønnow, T. F., Isakov, S. V., Wang, Z., Wecker, D., Lidar, D. A., Martinis, J. M., and Troyer, M. (2014). Evidence for quantum annealing with more than one hundred qubits. *Nature Physics*, 10(3):218–224.

- Boixo, S., Smelyanskiy, V. N., Shabani, A., Isakov, S. V., Dykman, M., Denchev, V. S., Amin, M. H., Smirnov, A. Y., Mohseni, M., and Neven, H. (2016). Computational multiqubit tunnelling in programmable quantum annealers. *Nature Communications*, 7(1):10327.
- Boros, E. and Gruber, A. (2014). On quadratization of pseudo-Boolean functions. arXiv:1404.6538 [math.OC].
- Boros, E. and Hammer, P. L. (2002). Pseudo-Boolean optimization. *Discrete Applied Mathematics*, 123(1):155–225.
- Childs, A. M., Farhi, E., and Preskill, J. (2001). Robustness of adiabatic quantum computation. *Physical Review A*, 65(1).
- Dattani, N. (2019). Quadratization in discrete optimization and quantum mechanics. arXiv:1901.04405 [quant-ph].
- Dattani, N. and Chancellor, N. (2019). Embedding quadratization gadgets on Chimera and Pegasus graphs. arXiv:1901.07676 [quant-ph].
- Dattani, N., Szalay, S., and Chancellor, N. (2019). Pegasus: The second connectivity graph for large-scale quantum annealing hardware. arXiv:1901.07636 [quant-ph].
- Denchev, V. S., Boixo, S., Isakov, S. V., Ding, N., Babbush, R., Smelyanskiy, V., Martinis, J., and Neven, H. (2016). What is the computational value of finite-range tunneling? *Phys. Rev. X*, 6:031015.
- Farhi, E., Goldstone, J., Gutmann, S., and Sipser, M. (2000). Quantum computation by adiabatic evolution. arXiv:0001106 [quant-ph].
- Fernández-Villaverde, J., Hurtado, S., and Nuño, G. (2023). Financial frictions and the wealth distribution. *Econometrica*, 91(3):869–901.
- Freedman, D. and Drineas, P. (2005). Energy minimization via graph cuts: settling what is possible. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 939–946 vol. 2.
- Glos, A., Kokainis, M., Mori, R., and Vihrovs, J. (2021). Quantum speedups for dynamic programming on n -dimensional lattice graphs. arXiv:2104.14384 [quant-ph].
- Grover, L. K. (1996). A fast quantum mechanical algorithm for database search. In Miller, G. L., editor, *Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, USA, May 22-24, 1996*, pages 212–219. ACM.

- Hull, I., Sattath, O., Diamanti, E., and Wendin, G. (2020). Quantum technology for economists. arXiv:2012.04473 [econ.GN].
- Ishikawa, H. (2014). Higher-order clique reduction without auxiliary variables. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1362–1369.
- King, J., Yarkoni, S., Nevisi, M. M., Hilton, J. P., and McGeoch, C. C. (2015). Benchmarking a quantum annealing processor with the time-to-target metric. arXiv:1508.05087 [quant-ph].
- Kolmogorov, V. and Zabini, R. (2004). What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159.
- Lucas, A. (2014). Ising formulations of many NP problems. *Frontiers in Physics*, 2.
- Messiah, A. (1958). *Quantum mechanics*. Dover Publications.
- Orús, R., Mugel, S., and Lizaso, E. (2019a). Forecasting financial crashes with quantum computing. *Physical Review A*, 99(6).
- Orús, R., Mugel, S., and Lizaso, E. (2019b). Quantum computing for finance: Overview and prospects. *Reviews in Physics*, 4.
- Pelofske, E., Hahn, G., and Djidjev, H. (2020). Advanced anneal paths for improved quantum annealing. arXiv:2009.05008 [quant-ph].
- Shin, S. W., Smith, G., Smolin, J. A., and Vazirani, U. (2014). How “quantum” is the D-Wave machine? arXiv:1401.7087 [quant-ph].
- Sweeting, A. (2013). Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry. *Econometrica*, 81(5):1763–1803.
- Tanburn, R., Okada, E., and Dattani, N. (2015). Reducing multi-qubit interactions in adiabatic quantum computation without adding auxiliary qubits. arXiv:1508.04816 [quant-ph].
- Taylor, J. B. and Uhlig, H. (1990). Solving nonlinear stochastic growth models: A comparison of alternative solution methods. *Journal of Business & Economic Statistics*, 8(1):1–17.
- Venegas-Andraca, S. E., Cruz-Santos, W., McGeoch, C., and Lanzagorta, M. (2018). A cross-disciplinary introduction to quantum annealing-based algorithms. *Contemporary Physics*, 59(2):174–197.
- Wootters, W. and Zurek, W. (1982). A single quantum cannot be cloned. *Nature*, 299(5886):802–803.

Zintchenko, I., Brown, E., and Troyer, M. (2015). Recent developments in quantum annealing.
Unpublished Manuscript.

Online Appendix

A Quadraticization Methods

We discuss four quadraticization methods that introduce either no auxiliary variables or the minimum number of auxiliary variables.

Deduction Reduction. Tanburn et al. (2015) propose a quadraticization strategy involving deducing properties that must hold in the ground state and performing variable substitutions accordingly. They demonstrate this on the simple constraint satisfaction problem:

$$\begin{aligned}x_1 + x_2 + x_3 &= 1 \\x_1x_4 + x_2x_5 &= x_3 \\x_1 + 2x_2 &= x_3 + 2x_4.\end{aligned}\tag{A.1}$$

The problem can be expressed in PBF form as:

$$\begin{aligned}\mathcal{H} &= (x_1 + x_2 + x_3 - 1)^2 + (x_1x_4 + x_2x_5 - x_3)^2 + (x_1 + 2x_2 - x_3 - 2x_4)^2 \\&= 2x_1x_2x_4x_5 - 2x_1x_3x_5 - 2x_2x_3x_5 - 2x_2x_3 + 6x_1x_2 - 3x_1x_4 - 8x_2x_4 + \\&\quad + x_2x_5 + 3x_2 + 4x_3x_4 + x_3 + 4x_4 + 1.\end{aligned}\tag{A.2}$$

For our purposes, it suffices to consider the first deduction-reduction. Equation (A.1) suggests that $x_1x_2 = x_2x_3 = x_3x_1 = 0$ in the ground state. Naive substitution of this deduction into Equation (A.2) reduces the Hamiltonian to the quadratic form:

$$\mathcal{H} = -3x_1x_4 - 8x_2x_4 + x_2x_5 + 3x_2 + 4x_3x_4 + x_3 + 4x_4 + 1.$$

Whereas the original Hamiltonian had a ground state energy level of $\mathcal{H} = 0$, the new Hamiltonian has states that yield $\mathcal{H} = -3$ and $\mathcal{H} = -2$. Thus, the substitution was not valid.

Tanburn et al. (2015) explain that this problem can be overcome by performing substitutions term-by-term and adding penalty terms where necessary. For instance, consider the quartic monomial term $2x_1x_2x_4x_5$. It can either add 0 or 2 to the Hamiltonian. In the ground state, $x_1x_2 = 0$, so it will add 0. However, it may add 2 or 0 outside of the ground state. Thus, imposing $x_1x_2 = 0$ will lower the energy level for some non-ground states. Thus, we must add $2x_1x_2$ as a penalty to \mathcal{H} if we perform this substitution. That is, $2x_1x_2x_4x_5 \rightarrow 2x_1x_2$.

Excludable Local Configurations (ELCs). Ishikawa (2014) defines an ELC as a partial assignment of variables that makes it impossible to achieve the ground state (global minimum). ELCs can be used to perform degree reduction without adding auxiliary variables. Dattani (2019) provides an example of a Hamiltonian that contains an ELC: $\mathcal{H} = x_1x_2 + x_2x_3 + x_3x_4 - 4x_1x_2x_3$.

Any state where $x_1x_2x_3 = 1$ will have a lower energy state than any state where $x_1x_2x_3 = 0$. So, the partial assignment $(x_1, x_2, x_3) = (1, 0, 0)$ is excludable, and we can apply the algorithm below, paraphrased from Ishikawa (2014), to reduce the cubic term:

1. If the coefficient on the higher-order term ζ is negative, find an ELC with a parity that matches the size of the partial assignment. If it is even, then find an ELC with the opposite parity.
2. For the ELC a , add the term to the Hamiltonian $\psi(x) = |\zeta| \prod_i \{a_i x_i + (1 - a_i)(1 - x_i)\}$.

In our case, $\zeta = -4$ and the partial assignment is $(x_1, x_2, x_3) = (1, 0, 0)$. This has an odd number of variables and an odd parity (number of 1s). Thus, condition 1 is satisfied. We can compute $\psi(x)$ as in:

$$\psi(x) = |-4|(1 * x_1 + 0 * (1 - x_1))(0 * x_2 + 1 * (1 - x_2))(0 * x_3 + 1 * (1 - x_3)).$$

Adding $\psi(x)$ to \mathcal{H} yields $\mathcal{H}' = x_1x_2 + x_2x_3 + x_3x_4 + 4x_1 - 4x_1x_2 - 4x_1x_3$, which has the same ground state as \mathcal{H} , but without the higher order terms.

Negative Term Reduction (NTR). When a higher-order term is negative, we can use the NTR approach (Kolmogorov and Zabin, 2004; Freedman and Drineas, 2005) to reduce it to a sum of quadratic terms using only one auxiliary variable, x_a :

$$-\prod_{i=1}^d x_i \rightarrow (d-1)x_a - \sum_{i=1}^d x_i x_a.$$

This form reproduces the full energy spectrum, including the ground state. It can also reduce d -order terms to quadratic terms using a single auxiliary variable.

Positive Term Reduction (PTR). For positive monomial terms, no algorithms can reduce d -order terms to quadratics using only one auxiliary variable. However, Boros and Gruber (2014) demonstrate how to reduce a d -order term to a quadratic using $d-2$ auxiliary variables:

$$\prod_{i=1}^d x_i \rightarrow \left(\sum_{i=1}^{d-2} x_{a_i} (d-i-1 + x_i - \sum_{j=i+1}^d x_j) \right) + x_{d-1} x_d.$$

B Annealing Schedule Simulations

We consider two trivial problems that require an iterative solution. The first is encoded in the Hamiltonian $\mathcal{H}_s = z_0 - z_1 - 2z_0z_1$. We set the initial state to $z_0 = z_1 = 0$ and apply Algorithm 5 with $\mathcal{C} = 1$. If z_0 is annealed first, it will be optimal to leave its value unchanged at 0, yielding $\mathcal{H}_s = 0$. And if z_1 is annealed next with $z_0 = 0$, it will be optimal to flip z_1 to 1, lowering \mathcal{H}_s to -1. In the absence of noise, an annealer would return this as the candidate solution, even though the global minimum is $\mathcal{H}_s = -3$ at $z_0 = z_1 = 1$. We must perform another iteration ($\mathcal{C} = 2$) to correct this. As shown in Table B.1, the $\mathcal{C} = 1$ case yields the correct answer in 10% of anneals, whereas $\mathcal{C} = 2$ boosts this success rate to 80%.

Cycles	Share Correct (\mathcal{H}_s)	Share Correct (\mathcal{H}_c)
1	0.10	0.00
2	0.80	1.00

Table B.1: Share of anneals that yield the correct solutions in the problems \mathcal{H}_s and \mathcal{H}_c given how many times the group 1 and group 2 annealing subschedules are repeated within a single anneal.

The second problem, which is more challenging and is encoded in the Hamiltonian $\mathcal{H}_c = z_2(2 + z_0 - 2z_0z_1) + z_3(2 - z_1 - 2z_0z_1)$, activates and de-activates components of the loss function using auxiliary qubits, z_2 and z_3 . This provides additional control over the implementation of cycles and aligns more closely with Algorithm 5. As with the previous problem, initializing in state $z_0 = z_1 = z_2 = z_3 = 0$ ensures that multiple cycles are needed to reliably obtain the global minimum. We show this experimentally in Table B.1, where $\mathcal{C} = 1$ yields the correct solution 0% of the time, but $\mathcal{C} = 2$ increases the success rate to 100%.