Title
Three Essays in Finance

Author
Fabisiak, Alex John

Publication Date
2019

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Three Essays in Finance

A dissertation submitted in partial satisfaction

of the requirements for the degree

Doctor of Philosophy in Management

by

Alex John Fabisiak

2019

ABSTRACT OF THE DISSERTATION

Three Essays in Finance

by

Alex John Fabisiak

Doctor of Philosophy in Management

University of California, Los Angeles, 2019

Professor Antonio E. Bernardo, Chair

In the first chapter, I apply machine learning techniques to numerically solve high-dimensional continuous time models in finance. Traditional methods rely on finite difference schemes for solutions to partial differential equations. By approximating the solution with a deep neural network, I am able to leverage the computational efficiency of neural networks and batch gradient descent to accurately compute solutions involving many state variables. I demonstrate the accuracy and efficiency of this method for Black-Scholes options pricing problems and dynamic programming problems in up to 50 spatial dimensions, far beyond the capability of grid methods. I also develop a solution method to mean field game type problems, where both a value function and a distribution function must solve a system of differential equations, utilizing mixture density networks.

In the second chapter (with Ivo Welch), we develop a model where buyers prefer local over lower-cost vendors even in the absence of direct preferences, taxes, subsidies, contracts, sanctions, information asymmetries, audits, etc. Instead, they prefer locals because they internalize the fact that local agents will in turn be more likely to buy from them in the future. Local sellers understand that buyers' preferences give them limited local market power, and therefore raise their prices and earn surplus in equilibrium.Our model can explain how voluntary reciprocity among subsets of *identical* agents can sustain itself, and how ex-ante *identical* goods from *ex-ante* identical sellers can acquire and maintain sustainably differentiated prices.

In the third chapter (with Antonio Bernardo and Ivo Welch), we develop a model where firms with lower leverage are not only less likely to experience financial distress but are also better positioned to acquire assets from other distressed firms. With endogenous asset sales and values, each firm's debt choice then depends on the choices of its industry peers. With indivisible assets, otherwise identical firms may adopt different debt policies—some choosing highly levered operations (to take advantage of ongoing debt benefits), others choosing more conservative policies to wait for acquisition opportunities. Our key empirical implication is that the acquisition channel can induce firms to *reduce* debt when assets become more redeployable. This article has been accepted for publication and is forthcoming in the *Journal of Financial and Quantitative Analysis*.

The dissertation of Alex John Fabisiak is approved.

Simon Adrian Board

William Giles Mann

Barney P Hartman-Glaser

Antonio E. Bernardo, Committee Chair

University of California, Los Angeles

2019

To Cara, for without her endless love and motivation,
I never could have made it this far.

And to my wonderful mom, dad, and brother,
for their support and encouragement from day one.

TABLE OF CONTENTS

1 Applications of Machine Learning to PDEs in Finance 1

1.1 Introduction 1

1.1.1 Literature 2

1.2 Methodology 4

1.2.1 Numerical PDE Methods 4

1.2.2 Machine Learning Principles 6

1.2.3 Procedure 11

1.3 Examples 13

1.3.1 Black-Scholes 13

1.3.2 Dynamic Programming and Hamilton-Jacobi-Bellman Equations . . . 19

1.3.3 Mean Field Games 22

1.4 Conclusion 27

1.5 Bibliography 37

2 Unenforced Endogenous Reciprocity 40

2.1 Introduction 40

2.2 The Model 42

2.2.1 Setup and Assumptions 42

2.2.2 Endogenous Quantities and Equilibrium Definition 45

2.2.3 Illustration 48

2.3 Benchmark: Non-Reciprocity 49

2.4 Full Voluntary Reciprocity 50

2.4.1 Comparing Equilibria 52

2.5 Partial Voluntary Reciprocity . . . 53
2.6 Equilibrium Comparisons . . . 55
2.7 Discussion . . . 56
2.7.1 Period-by-Period Seller Optimization . . . 56
2.7.2 Alternative Pricing Mechanisms . . . 57
2.7.3 Alternative Arrangements and Larger Networks . . . 58
2.7.4 Assumptions—Tractable vs Economic and Extensions . . . 59
2.8 Conclusion . . . 61
2.9 Appendix: Equilibrium Derivation . . . 62
2.9.1 Sellers' Problem . . . 63
2.9.2 Wealth Dynamics . . . 64
2.9.3 Buyer's Problem . . . 65
2.10 Appendix: Stage Game Without Premiums . . . 70
2.11 Appendix: Stage Game With Premiums . . . 73
2.12 Bibliography . . . 85

3 Asset Redeployability, Liquidation Value, and Endogenous Capital Structure Heterogeneity . . . 86
3.1 Introduction . . . 86
3.2 The Opportunistic-Acquisition Channel . . . 89
3.2.1 Model Setup and Assumptions . . . 89
3.2.2 Equilibrium . . . 95
3.2.3 Solution . . . 95
3.2.4 Implications . . . 97
3.3 The Distress-Reorganization Channel . . . 100

3.3.1 Setup and Assumptions 101
3.3.2 Access to Infinite Financing / Eliminating The Acquisition Channel . 104
3.3.3 Equilibrium With Acquisitions and Financial Distress Channels . . . 105
3.3.4 Implications 108
3.4 Discussion and Literature Context 115
3.4.1 Welfare 115
3.4.2 Generalizations of the Model 117
3.4.3 Related Literature 119
3.4.4 Relation to Shleifer-Vishny 1992 120
3.4.5 Summary 123
3.5 Conclusion 123
3.6 Appendix: Proof of Theorem 1 124
3.6.1 Pure-strategy interior equilibrium 124
3.6.2 Mixed-strategy equilibrium 125
3.6.3 Pure-strategy extreme equilibrium 126
3.6.4 Uniqueness 126
3.7 Appendix: Extension of Theorem 2 to the Full Parameter Space 127
3.7.1 Statement 127
3.7.2 Proof 130
3.8 Appendix: Extension to Immediate Use of Debt Benefits 143
3.9 Appendix: Extension to Industry Booms and Busts 144
3.9.1 Reallocation 144
3.9.2 Pro-Cyclical 144
3.9.3 Counter-Cyclical 145
3.10 Appendix: Tax Shields as Sources of Debt Benefits 145

3.11 Bibliography 161

LIST OF FIGURES

1.1 Fixed sampling may ignore problem points 29

1.2 Error for the European Call Option 30

1.3 Absolute Error of Neural Network on Black-Scholes with Variable Parameters ($r = 10\%, \sigma = 20\%, K = 10$) 31

1.4 Error for 1-Dimensional HJB Equation (Noise from Monte-Carlo Reference) . . 32

1.5 1-Dimensional Mean Field Game 33

1.6 2-Dimensional Mean Field Game 34

2.1 Time Line 44

2.2 Mutual Best Response 77

2.3 A's Payoff as Function of His Imposed Handicap, Given B's Handicap 78

2.4 Game Dynamics: Winning Probabilities and Pricing as Function of Handicap in Single Period 79

2.5 Non-Equilibrium Profits in A Single Period 80

2.6 Comparative Statics for Equilibrium Handicaps 81

2.7 Comparative Statics for Equilibrium Wealth 82

2.8 Comparative Statics for Neighbor to Neighbor Prices Paid 83

2.9 Comparative Statics for Probability of Neighbor Purchases 84

3.1 Game Tree for the Acquisition Model 148

3.2 Comparative Statics for Heterogeneity in the Acquisition-Only Model ($\phi = 0$) . 149

3.3 Comparative Statics for the Acquisition-Only Model Leverage (ϕ=0) 150

3.4 Peer Effects on Debt Choice 151

3.5 Game Tree for the Full Model ($F_i > P$) 152

3.6 Comparative Statics for Heterogeneity when ϕ=0.25 153

3.7 Comparative Statics for Industry Leverage when ϕ=0.25 154

3.8 Ancillary Comparative Statics for ϕ=0.25 155

3.9 Allocational Efficiency 156

3.10 Value and Debt-Ratio When Benefits are Tax Shields 157

LIST OF TABLES

1.1 Performance of Neural Network on Multi-dimensional Black-Scholes 29

1.2 Performance of Mathematica `NDSolve` on Multi-dimensional Black-Scholes . . . 31

1.3 Performance of Neural Network on Black-Scholes with Variable Parameters . . . 31

1.4 Performance of Neural Network Using Only t on Multi-dimensional Black-Scholes 35

1.5 Performance of Neural Network on Multi-dimensional HJB 35

1.6 Performance of Mathematica `NDSolve` on Multi-dimensional HJB 36

3.1 Variables 158

3.2 Summary of Comparative Statics 159

3.3 Model Implication and Features Comparison 160

ACKNOWLEDGMENTS

I am deeply grateful to my advisor, Tony Bernardo, for his invaluable guidance and support throughout my time at Anderson. I am also indebted to Ivo Welch, who has been a terrific mentor and co-author. Additionally, I am very thankful to Stavros Panageas, Barney Hartman-Glaser, William Mann, and Simon Board for their advice and support.

I would like to thank my fellow Ph.D. students, with special mention to Darren Aiello, Chady Gemayel, Mahyar Kargar, and Geoff Zheng. They have all become lifelong friends, and their advice, feedback, and encouragement have been essential to this journey.

The administrative staff, Craig Jessen, Kristin Christian, Lydia Heyman, and Leanna Cortez, have been indispensable in helping me navigate this program. I would also like to thank the Anderson School of Management and the Fink Center, for their generous financial assistance.

VITA

2008–2012	B.S. Business Administration (Finance), University of Colorado–Boulder
2008–2012	B.A. Mathematics, University of Colorado–Boulder
2010–2012	Intern–Financial Analyst, Sargent Bickham Lagudis Investment Management
2011–2012	Teaching Assistant, University of Colorado–Boulder. Taught Introduction to Finance
2011	Intern–Credit Analyst, Silver Lake Credit
2012	Research Assistant, University of Colorado–Boulder
2012	Research Assistant, University of Colorado–Boulder
2013–2019	Ph.D. Candidate in Management (Finance), UCLA Anderson School of Management
2013–2017	Fink Center Scholar
2013–2017	Anderson Fellow
2014–2018	Teaching Assistant, UCLA Anderson School of Management. Taught Stochastic Calculus, Corporate Finance, Prescriptive Models and Data Analytics, Foundations of Finance
2018	Instructor, UCLA Anderson School of Management. Taught Introduction to Python

CHAPTER 1

Applications of Machine Learning to PDEs in Finance

1.1 Introduction

Continuous time models have become a critical tool within many fields of modern finance. They were first applied to asset pricing by Merton (1973) and Black and Scholes (1973). Since then, they have been adapted to a range of other topics: Cox et al. (1985) for term structure models, Leland (1994) for corporate finance, and Sannikov (2008) for contracting. By taking models to the continuous time limit, researchers are able to model complex interactions without sacrificing tractability or intuition.

The solution at the heart of these continuous time models can ultimately be characterized by a partial differential equation (PDE). In a number of cases, these have analytic global solutions that describe the model's evolution. Unfortunately, analytic solutions are rare, so numerical approximation methods are required. Simple numerical methods, such as finite differences, are often highly effective when the PDE is linear or is only a function of time and a single state variable. For more complex problems, more complex techniques are required.

In this paper, we describe a numerical PDE solution method based on machine learning techniques. This method is well suited to non-linear or high dimensional problems with a large number of heterogeneous agents or economic state variables. By modeling the underlying approximating function as the output of a deep neural network (DNN), we are able to employ many of the latest methodologies and highly optimized software packages from the machine learning field.

Recent advances in machine learning have focused on the problem of efficiently estimating the high-dimensional, non-linear relationships which underlie observational data. To this end, a DNN can be thought of as a flexible, yet computationally efficient, approximating function. The objective is generally to minimize a least squares loss, where the methods look to minimize the distance between the approximation and observed data.

Analogously, our method tries to minimize the squared residuals of the PDE, which measure how well the approximation satisfies the required PDE. We are then able to utilize the backpropagation algorithm for efficient derivative computation, and batch gradient descent to evaluate the loss on only a randomly sampled meshless subset of the domain. These techniques make the solution method well suited to non-linear and high-dimensional problems.

Through examples with known analytic solutions, we demonstrate the effectiveness of our method. When applied to pricing multi-asset basket options with Black-Scholes, our method surpasses the performance of finite difference methods beyond three assets. Beyond six assets, finite difference methods are infeasible due to memory constraints, while the neural network based approach is able to handle up to 50 assets with error under 0.3%. Similarly low error is achieved when our method is applied to a non-linear dynamic control problem.

This paper also develops a novel application of mixture density networks to PDEs of density functions, such as the Kolmogorov-forward equation. This neural network extension models the probability density function of a mixture of normal random variables, where mean, variance, and mixing probabilities change over time. The resulting PDF is highly flexible and capable of approximating many other density functions. We demonstrate their usefulness by solving a mean field game, which involves a coupled Hamilton-Jacobi-Bellman equation and a Kolmogorov-forward equation.

1.1.1 Literature

Because dynamic models are ubiquitous in finance and economics, methods for solving more complex models are an area of active research. In discrete time, Fernandez-Villaverde, Rubio-Ramirez, and Schorfheide (2016) describes the perturbation and projection techniques. More recent work has been concerned with using sparse representations to increase the number of possible dimensions (Winschel and Kratzig (2010), Brumm and Scheidegger (2017)).

In continuous time, solution methods generally involve the use of finite difference schemes, which approximate the derivatives as differences on a grid. In some cases, it is possible to slightly reduce the dimensionality of these problems using the shooting method or the method of lines (Brunnermeier and Sannikov (2016)). However, they are still subject to the "curse of dimensionality," where the memory requirements grow exponentially in the number of dimensions.

Outside the economics literature, neural networks were first applied to PDEs by Dissanayake and Phan-Thien (1994). Kumar and Yadav (2011) surveys the many variations that have been developed. However, only very recently have they been applied to high-dimension problems, as in Han, Jentzen, and E (2017), which produces solutions at a single point.

This paper's contribution is to be among the first to bring techniques from machine learning to continuous time models in finance. One other such proposal is Duarte (2018), which embeds neural networks in traditional fixed point iteration. A limitation of this approach is the reliance on the contraction mapping theorem for convergence.

Sirignano and Spiliopoulos (2017) develops a very similar methodology to ours, which involves batching and minimizing PDE residuals. However, their method uses a LSTM style neural network, which is more complex and slower to train. We achieve similar accuracy using a simpler DNN, and explore the importance of sampling methods on performance.

Finally, our paper proposes a novel machine learning based approach to solve mean field game problems. Mean field games describe the evolution of a continuum of agents interacting and acting optimally (Lasry and Lions (2007)). These problems are characterized by a coupled system of PDEs, where one solution is a probability density function. Current solution methods are highly customized and rely on finite difference methods (Achdou, Han, et al. (2017)). Our method proposes an alternative solution technique based on mixture density networks (Bishop (1994)).

1.2 Methodology

1.2.1 Numerical PDE Methods

Consider a generic differential equation of the form

$$\mathcal{L}F(\vec{x}) = g(\vec{x}), \quad \vec{x} \in D \tag{1.1}$$

$$\mathcal{B}F(\vec{x}) = h(\vec{x}), \quad \vec{x} \in B \tag{1.2}$$

where $\mathcal{L}$ and $\mathcal{B}$ are some differential operators. A strong solution $F(\cdot)$ would exactly satisfy both the PDE, (1.1), over the entire domain D, as well as the boundary conditions, (1.2), over the boundary B. Since this is rarely possible analytically, many numerical approximation methods have been developed.

The method in this paper falls in the class of residual methods, which include finite elements, collocation, and spectral methods (Judd (1998), Boyd (2001)). First, they construct an approximating function $f(x;\theta)$ that built on some basis function, where θ are the underlying parameters. These functions could be a high order polynomials as in Chebyshev methods, sine functions in spectral methods, or local piecewise functions in finite-element methods. Then, select some finite collection of points $\hat{D} \subset D$ and $\hat{B} \subset B$, generally on a regularly spaced grid. At these points, the residuals are computed:

$$R_1(\vec{x}_i, f(\vec{x}_i;\theta)) = \mathcal{L}f(\vec{x}_i;\theta) - g(\vec{x}_i), \quad \vec{x}_i \in \hat{D}$$

$$R_2(\vec{x}_i, f(\vec{x}_i;\theta)) = \mathcal{B}f(\vec{x}_i;\theta) - h(\vec{x}_i), \quad \vec{x}_i \in \hat{B}$$

In cases where $f(x;\theta)$ is sufficiently flexible and θ has enough degrees of freedom, it may be possible to jointly set all the residuals to zero. Much work has been done to formulate such problems as matrix inversion problems that can be exactly solved, such as with linear Chebyshev methods. However, these methods are often restrictive in the particular PDE forms that can be accommodated.

Our method differs from these traditional methods in several important ways. First, we will construct $f(\vec{x};\theta)$ as the output of a dense, feedforward neural network. Neural networks are well suited to this task, as will be discussed shortly. Second, our goal will be to minimize

$$L(\theta;\hat{D},\hat{B}) = \sum_{\vec{x}_i \in \hat{D}} R_1(\vec{x}_i, f(\vec{x}_i;\theta))^2 + \lambda \sum_{\vec{x}_i \in \hat{B}} R_2(\vec{x}_i, f(\vec{x}_i;\theta))^2$$

where $\lambda > 0$ controls the relative weighting of the PDE versus the boundary residuals. Instead of exactly solving to set all residuals to zero, we will instead be minimizing the least squares loss.

Both of these points are common within the existing literature on neural network PDE methods. However, some papers will reformulate the problem so that the boundary conditions are always satisfied (Lagaris, Likas, and Fotiadis (1998)). For example, with initial condition $F(0,x) = g(x)$, they construct $f(t,x;\theta) = g(x) + t \cdot \tilde{f}(t,x;\theta)$ where $\tilde{f}$ is the neural network output. This way, the initial condition is always trivially satisfied. We do not take this approach here, as performance was found to be extremely poor in practice when g was not differentiable, like in Black-Scholes.

Finally, our method will find the minimizing parameter values, θ^*, using Batch Gradient Descent (BGD) and the Adam Optimizer. Instead of constructing $\hat{D}$ as a fixed grid, we will instead draw random samples repeatedly from the space D. This makes our method a meshless method that scales better to high dimensions. Methods that rely on a fixed grid of points are bound by the "curse of dimensionality," as the number of required points grows exponentially. Only recently has BGD been utilized for solving PDEs, as in Sirignano and Spiliopoulos (2017) and Duarte (2018). By resampling repeatedly, the goal is to improve the fit of the neural network over the entire domain D, and not at a sparse few points. The Adam Optimizer is helpful in this context, as it exhibits momentum which will retain a memory of previous points we have sampled.

1.2.2 Machine Learning Principles

1.2.2.1 Artificial Neural Networks

An artificial neural network (ANN) is a function from $f : \mathbb{R}^d \to \mathbb{R}$ that is composed of many simpler functions, or neurons. We can think of a single one of these neurons much like a logistic regression. Given some input point $\vec{x} \in \mathbb{R}^d$, a single neuron would take the value $h(\vec{x}) = g(\vec{w} \cdot \vec{x} + b)$ where $g(z)$ is the logistic function $g(z) = \frac{1}{1+e^{-z}}$.

Clearly, this logistic regression output $h(\vec{x})$ is not sufficiently flexible to solve most PDEs. The free parameters we would have to choose are $\vec{w}$ and b, but the output is bounded in $(0, 1)$ and can only model functions that are monotone in $\vec{x}$. Instead, the ANN uses the value of many such neurons as intermediate steps. Let

$$\vec{h}(\vec{x}) = g(W_1 \cdot \vec{x} + \vec{b}_1)$$

where $\vec{h}, \vec{b} \in \mathbb{R}^m$ and $W_1 \in \mathbb{R}^{m \times d}$. We can think of this as the stacked output of m logistic regressions, and is known as the hidden layer. The output returned by the ANN is then

$$f(\vec{x}) = W_2 \cdot \vec{h}(\vec{x}) + b_2$$

where $W_2 \in \mathbb{R}^{1\times m}$ and b_2 is a scalar.

This function can be computed efficiently, as it is mostly linear operations with the only non-linearity being the logistic function $g(\cdot)$. While we have so far used the logistic function, there are other potential activation functions $g(\cdot)$. Popular functions include the hyperbolic tangent function $g(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$, the rectified linear function $g(z) = \max(0, z)$, or the leaky rectified linear function $g(z) = \max(\alpha z, z)$. It is worth noting that the rectified linear functions are not suited to PDE approximation, as the resulting output is piecewise linear. Therefore, the second derivatives with respect to $\vec{x}$ are almost always zero.

Despite the computational simplicity of ANNs, they have also proven extremely flexible. The *Universal Approximation Theorem* (Cybenko (1989)) shows that such an ANN is capable of approximating any Borel measurable function arbitrarily well, given enough hidden neurons, m. And, of particular use for PDE applications, the function is also able to approximate the derivatives of any such function arbitrarily well (Hornik (1991)).

Unfortunately, the number of hidden neurons required for a good approximation may be intractably large. This can be mitigated by constructing a deep neural network (DNN). Instead of only having one layer of hidden neurons, we construct multiple hidden layers as functions of the preceding hidden neurons.

$$
\begin{aligned}
\vec{z}_1 &= W_1 \cdot \vec{x} + \vec{b}_1 \\
\vec{h}_1 &= g(\vec{z}_1) \\
\vec{z}_2 &= W_2 \cdot \vec{h}_1 + \vec{b}_2 \\
\vec{h}_2 &= g(\vec{z}_2) \\
&\vdots \\
\vec{z}_l &= W_l \cdot \vec{h}_{l-1} + \vec{b}_l \\
\vec{h}_l &= g(\vec{z}_l) \\
f &= W_{l+1} \cdot \vec{h}_l + b_{l+1}
\end{aligned}
\tag{1.3}
$$

As the number of hidden layers increases, the DNN is capable of approximating more complex functions with fewer hidden neurons in each layer and fewer total parameters.

1.2.2.2 Gradients

Part of the reason for the ubiquitous use of DNNs, is the ease with which gradients are computed. As part of the optimization process, algorithms require computation of the derivatives of the loss function with respect to the parameters. Here however, we are interested in the derivatives with respect to the input dimension. Consider below that x is scalar, and there is only one hidden neuron at each layer, so h_i is also scalar.

$$
\begin{aligned}
\frac{\partial f}{\partial x} &= \frac{\partial f}{\partial h_l} \cdot \frac{\partial h_l}{\partial z_l} \cdots \frac{\partial h_1}{\partial z_1} \cdot \frac{\partial z_1}{\partial x} \\
&= W_{l+1} \cdot g'(z_l) \cdot W_l \cdots g'(z_1) \cdot W_1 \\
&= W_{l+1} \cdot h_l(1 - h_l) \cdot W_l \cdots h_1(1 - h_1) \cdot W_1
\end{aligned}
$$

If we have chosen $g(z)$ to be the logistic function, then conveniently $g'(z) = g(z)(1-g(z))$.

The above derivative is just the product of many terms we already computed. On the forward pass, where we compute f, we store all the intermediate values along the way. Then, to calculate the derivative (the backward pass), we are able to simply multiply values we have already stored. This algorithm is known as the backpropagation algorithm. For a more detailed explanation of this algorithm, and neural networks in general, see Goodfellow, Bengio, and Courville (2016). The higher order derivatives become somewhat more complicated, as each h_j depends on x as well, and the product rule applies. However, they remain a linear combination of previously computed quantities.

Because this process is so essential to machine learning, it is handled entirely within the implementations of popular machine learning toolboxes (TensorFlow, Keras, Matlab, etc.). Also, it is entirely linear algebra based so can be computed extremely quickly with Graphical Processing Units (GPUs). This means that for researchers, solving PDE problems in these toolboxes is as simple as writing down the differential equation. All of the computation of the derivatives (and gradients with respect to parameters), it taken care of automatically.

1.2.2.3 Batch Gradient Descent

Another essential component of modern machine learning techniques is Batch Gradient Descent, a relative of Stochastic Gradient Descent (Robbins and Monro (1951)). Suppose we have 10,000 fixed points in D that we wish to fit the PDE at, and suppose the boundary condition is always satisfied. Then compute the loss for a given parameter estimate as:

$$L(\theta_i; D) = \frac{1}{10,000} \sum_{j=1}^{10,000} R_1(\vec{x}_j, f(\vec{x}_j; \theta_i))^2$$

To improve our fit, we would like to take a step in the direction of decreasing gradient.

$$\begin{aligned} \theta_{i+1} &= \theta_i - \gamma \nabla_\theta L(\theta_i; D) \\ &= \theta_i - \gamma \frac{1}{10,000} \sum_{j=1}^{10,000} \nabla_\theta R_1(\vec{x}_j, f(\vec{x}_j; \theta_i))^2 \end{aligned}$$

Here γ is the step size, also known as the learning rate.

The gradient above resembles the population mean. If we assign each point in D equal probability and draw a subsample of n points, $\hat{D} \subset D$, then we can instead update our parameters using the sample mean:

$$\theta_{i+1} = \theta_i - \gamma \frac{1}{n} \sum_{x_j \in \hat{D}} \nabla_\theta R_1(\vec{x}_j, f(\vec{x}_j; \theta_i))^2$$

This sample mean will be equal to the true gradient in expectation, and the standard deviation decreases as $1/\sqrt{n}$. This subsample $\hat{D}$ is what is known as a batch. By using a smaller sample, we are able to update θ_i in a direction that is correct on average. The case when $n = 1$ is known as Stochastic Gradient Descent. While sampling greater n provides a less noisy estimate, the returns are decreasing. In order to double the precision, we need to compute the residuals for four times as many points. Ultimately, the tradeoff between speed and accuracy is problem specific.

Another important point comes up when using machine learning for PDEs. Unlike with most machine learning contexts, we have access to unlimited training data. We can sample any point in D as many times as we like. However, almost all prior methods restrict attention to a fixed grid of points within D.

Figure 1.1 illustrates a potential problem with this method. Because there are a multitude of solutions to any PDE with different boundary conditions, it is possible that our approximating function jumps between them. We may satisfy the boundary condition exactly, and the residual of the PDE may be zero at the grid points. However, because the PDE is not satisfied in between, we have moved to a solution with the wrong boundary condition. To remedy this, we instead draw different points from D for every step of our gradient descent. This way, the problem region is hopefully sampled and remedied.

1.2.2.4 Adam Optimizer

The Adam Optimizer is an extremely popular and effective extension of SGD (Kingma and Ba (2014)). Because the gradient estimates are noisy, this method estimates the current gradient as an AR(1) process. This feature is called momentum and is a popular component of many optimizers. It allows the optimizer to retain a memory of the past points that have been sampled. Even though we only sample a few points at every iteration, this memory tries to eliminate the possibility of backtracking and undoing previous updates.

The second moment of the gradient is also estimated as an AR(1). Then, the step size is inversely proportional to the variance of the gradient. When the gradient estimates are noisier, the algorithm will take smaller steps.

$$\begin{aligned}
g_i &= \frac{1}{n}\sum_{x_j \in \hat{D}} \nabla_\theta R_1(\vec{x}_j, f(\vec{x}_j;\theta_i))^2 \\
m_i &= \beta_1 m_{i-1} + (1-\beta_1) g_i \\
v_i &= \beta_2 v_{i-1} + (1-\beta_2) g_i^2 \\
\gamma_i &= \gamma \frac{\sqrt{1-\beta_2^i}}{1-\beta_1^i} \\
\theta_{i+1} &= \theta_i - \gamma_i \frac{m_i}{\sqrt{v_i}+\epsilon}
\end{aligned} \tag{1.4}$$

Equation (1.4) describes the process in detail. g_i is the gradient estimate from the current batch. m_i and v_i are the AR(1) estimates of the gradient and its second moment. The learning rate γ_i is defined to eliminate the bias of initializing the process at 0. Finally, the last line describes the parameter update. The authors describe $m_i/\sqrt{v_i}$ as a signal to noise ratio, where the step size decreasing when the SNR is low.

1.2.2.5 Mixture Density Networks

A useful variation of the DNN described above is the mixture density network (Bishop (1994)). Often in economic models, researchers are interested in the evolution of distribution functions over time. Achdou, Buera, et al. (2014) reviews several such continuous time models, where the evolution of agent distributions is governed by a differential equation, known as the Kolmogorov Forward Equation.

As before, we wish to learn an approximating function $d(t, x; \theta)$ that satisfies some PDE. (Let us assume $x \in D \subset \mathbb{R}$ for notational simplicity.) However, we also require that this function is a valid PDF so satisfies the constraint that $\int_D d(t, x; \theta)dx = 1$. One way to do this would be to use an DNN like above, but to normalize its value:

$$d(t, x; \theta) = \frac{f(t, x; \theta)}{\int_D f(t, x; \theta) dx}.$$

Computing this integral is extremely costly numerically.

Instead, we will construct the function as the PDF of a mixture of normals, where the mean, variance, and mixing probabilities vary over time. To begin, we will compute a hidden layer in the same method as above. However, the only input variable will be the time index t:

$$\vec{h}_1(t) = g(W_1 \cdot t + \vec{b}_1)$$

$$\vec{z}_2 = W_2 \cdot \vec{h}_1 + \vec{b}_2$$

We will then partition $\vec{z}_2$ into three parts. $\vec{z}_2 = [\vec{\mu}|\vec{\sigma}|\vec{p}]$, where $\vec{\mu}, \vec{\sigma}, \vec{p} \in \mathbb{R}^k$, and k is the number of normal distributions we mix between. We transform $\vec{\sigma} = e^{\vec{\sigma}}$ so it is always positive, and $p_i = \frac{e^{-p_i}}{\sum_j e^{-p_j}}$ with the softmax function so that they are discrete probabilities.

Then define:

$$d(t, x; \theta) = \prod_{i=1}^{k} p_i \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \frac{(x - \mu_i)^2}{2\sigma_i^2} \tag{1.5}$$

Now it is guaranteed that $d(t, x; \theta)$ will be a valid PDF for any choice of θ (and thereby $\vec{h}$).

Another possible route would be to define a CDF that is monotone with range $[0, 1]$. This has not been explored in this paper, but some work on monotone neural networks has been done (You et al. (2017)). This method is also similar to the work on Radial Basis Function Networks that have been used as universal approximators for solving PDEs (Kumar and Yadav (2011)), but here we are able to normalize them with known constants.

1.2.3 Procedure

For the generic PDE:

$$\mathcal{L}F(\vec{x}) = g(\vec{x}), \quad \vec{x} \in D$$

$$\mathcal{B}F(\vec{x}) = h(\vec{x}), \quad \vec{x} \in B$$

the procedures for this method are as follows:

0. $i = 0$. Initialize our approximating dnn $f(\vec{x}; \theta_0)$ with random parameters θ_0.

1. Draw sample of interior points $\vec{x}_j \in D$. Call this batch $\hat{D}_i$.

2. Compute the PDE residuals at these points.

$$R_1(\vec{x}_j, f(\vec{x}_j; \theta_i)) = \mathcal{L} f(\vec{x}_j; \theta_i) - g(\vec{x}_j), \quad \vec{x}_j \in \hat{D}_i$$

3. Draw sample of boundary points $\vec{x}_j \in B$. Call this batch $\hat{B}_i$.

4. Compute the boundary residuals at these points.

$$R_2(\vec{x}_j, f(\vec{x}_j; \theta_i)) = \mathcal{B} f(\vec{x}_j; \theta_i) - h(\vec{x}_j), \quad \vec{x}_j \in \hat{B}_i$$

5. Compute the total loss.

$$L(\theta_i; \hat{D}_i, \hat{B}_i) = \sum_{\vec{x}_j \in \hat{D}_i} R_1(\vec{x}_j, f(\vec{x}_j; \theta_i))^2 + \lambda \sum_{\vec{x}_j \in \hat{B}_i} R_2(\vec{x}_j, f(\vec{x}_j; \theta_i))^2$$

6. If the total loss is sufficiently small, end. Otherwise, update the parameters of our approximating function (assign θ_{i+1}) using the Adam optimizer to minimize the total loss. Update the learning rate, set $i = i + 1$, and return to step 1.

Several choices must be made by the researcher when using this method. Many of these choices are common to all machine learning methods: the complexity of the neural network (number of layers and hidden neurons), the number of samples per batch, the initial learning rate and possible learning rate decay, the number of iterations, and optimizer parameters. The selection of these parameters is known as hyperparameter tuning, and there are many resources online about rule-of-thumb methods, or more rigorous Bayesian methods (Snoek, Larochelle, and Adams (2012)).

The question is, how do we know which parameter values produce a better solution to our PDE? In cases where we do not know the analytic solution, this may be difficult to judge. One important way to evaluate the fit is to check that the solution matches the known boundary conditions well. Running through the method with the loss only being equal to the squared boundary residuals should produce a good approximation on the boundary. A bad fit may be indicative of bad optimizer parameters, or that the chosen neural network structure is insufficiently flexible to model the solution.

The value of the loss function at a fresh sample of evaluation points is also a good measure of fit. Because we are able to compute the exact PDE residuals at any point in D, there is no concern of overfitting. That is, with observational data, the function may improve the loss by catering to every outlier, at the expense of generalizability. Here, there is a strict correlation between the current loss value and the quality of the approximation in general.

Another critical choice is how to draw the sample points. The main concern is that the procedure will focus most on fitting the region of the space where the most points are drawn. The loss function will weight different regions in the proportion with which they are sampled, so comparing loss measures between sampling methods is invalid. A method which only ever samples the same point will be able to achieve loss close to zero, but it would not be considered a good approximation. The subsequent examples illustrate some potential methods and the tensions they create.

1.3 Examples

1.3.1 Black-Scholes

1.3.1.1 One Dimension

We begin with a simple European call option for illustration. Suppose we have a risk-free asset that earns rate r, and a stock with dynamics

$$\frac{dS(t)}{S(t)} = \mu dt + \sigma dW(t)$$

At time T, the option will have a terminal payoff $\Phi(S(T)) = \max\{S(T) - K, 0\}$. We know that the price of this call option satisfies the Black-Scholes PDE:

$$\begin{aligned} &\frac{\partial F(t,S)}{\partial t} + rS\frac{\partial F(t,S)}{\partial S} + \frac{1}{2}S^2\sigma^2\frac{\partial^2 F(t,S)}{\partial S^2} - rF(t,S) = 0 \quad && \forall (t,S) \in [0,T] \times [0,\infty) \\ \text{s.t.} \quad &F(T,S) = \Phi(S) && \forall S \in [0,\infty) \end{aligned}$$

While the analytic solution will solve the PDE exactly on the whole domain, what we seek is an approximating function that gets close. In particular, we will will use $f(t, S; \theta)$ that is the output of a deep neural-network with four hidden layers of 50 neurons each, using the tanh activation function. Also, we cannot sample $S \in [0, \infty)$, so we will instead sample $S \in [0, \overline{S}]$.

0. $i = 0$. Initialize our approximating function as a DNN with four hidden layers of 50 neurons each: $f(t, S; \theta_0)$ with random parameters θ_0.

1. Draw sample of interior points $\hat{D}_i$. Sample $(t_j, S_j) \in [0, T] \times [0, \overline{S}]$ uniformly.

2. Compute the PDE residuals at these points.

$$R_1(t_j, S_j, f(t_j, S_j; \theta_i)) = \frac{\partial f(t_j, S_j; \theta_i)}{\partial t} + rS_j\frac{\partial f(t_j, S_j; \theta_i)}{\partial S} + \frac{1}{2}S_j^2\sigma^2\frac{\partial^2 f(t_j, S_j; \theta_i)}{\partial S_j^2} - rf(t_j, S_j; \theta_i)$$

3. Draw sample of boundary points $\hat{B}_i$. Sample $(t_j, S_j) \in [T, T] \times [0, \overline{S}]$ uniformly.

4. Compute the boundary residuals at these points.

$$R_2(t_j, S_j, f(t_j, S_j; \theta_i)) = f(T, S_j; \theta_i) - \max\{S_j - K, 0\}$$

5. Compute the total loss.

$$L(\theta_i; \hat{D}_i, \hat{B}_i) = \sum_{(t_j,S_j)\in\hat{D}_i} R_1(t_j, S_j, f(t_j, S_j; \theta_i))^2 + \lambda \sum_{(t_j,S_j)\in\hat{B}_i} R_2(t_j, S_j, f(t_j, S_j; \theta_i))^2$$

6. If the total loss is sufficiently small, end. Otherwise, update the parameters of our approximating function (assign θ_{i+1}) using the ADAM optimizer to minimize the total loss. Update the learning rate, set $i = i + 1$, and return to step 1.

For this example, let $r = 5\%$, $\sigma = 20\%$, $T = 5$, and $K = 10$ and sample share prices on the range $[0, 50]$. We make 40,000 iterations, each time sampling 1,000 points on the interior and on the boundary. We start with a learning rate of 0.01 and decrease it by 10% every 500 iterations. All of our examples were coded in Google's open source TensorFlow package for Python.

Figure 1.2 illustrates the quality of our fit. Figure 1.2 (a) shows the error in the fit over the entire rectangular region. We see that our method does quite poorly in the top left corner. This is due to the fact that our boundary condition only holds for the right edge (at maturity), and for a limited region ([0,50]). If we ask for the value of the option at point A, (t,S)=(1,40), the value at this point depends strongly on what the payoff will be for values of $S > 50$, as it is relatively likely that the final share price will fall there at maturity. Because this is outside of our selected region, the value in this corner is largely undetermined.

In Figure 1.2 (c), we use the same approximation, but restrict our attention to points where the terminal share price is unlikely to fall outside our sampled region. For points within the displayed region, it would take greater than a three standard deviation return to make it outside the modeled boundary region. The magnitude of error is much smaller in this region, with the biggest error being at the strike price at maturity. This is a result of trying to match a non-smooth function with a smooth one.

Another possible method is to sample $\hat{D}_i$ and $\hat{B}_i$ according to the risk-neutral diffusion process. For the interior points, sample $t \in [0, T]$ uniformly, and sample $S = S_0 e^{(r - \sigma^2/2)t + \sigma\sqrt{t}Z}$ where Z is standard normal. For the boundary points, fix $t = T$ and sample S as before.

Figure 1.2 also illustrates the performance this GBM sampling method, given a starting price of $S_0 = 10$. We see in Figure 1.2 (b) that the approximation is still poor in the top left, and now, bottom left corners. However, we have also not sampled many points in this region. As discussed above, sampling points in this region would be wasteful without sampling from the boundary condition in the corresponding region. Figure 1.2 (d) shows the same approximation, but subset to the region of returns less than three standard deviations from S_0. The worst fit is near the edges, which is due to insufficient sampling of their corresponding boundary region.

Sampling test points according to the GBM, the uniform sampling method produces an average pricing error of \$0.0036 compared to \$0.0015 for GBM sampling. The percent error at S_0 is 0.095% with uniform sampling, compared to 0.002% for GBM sampling. This is due to the fact that the GBM method does not waste time fitting the approximation in regions that are not defined.

1.3.1.2 Multiple Assets

Now we will apply our method to the multi-dimensional Black-Scholes PDE. Suppose we have a risk-free asset the earns rate r, and n stocks with dynamics

$$\frac{dS_i(t)}{S_i(t)} = \mu_i dt + \sigma_i dW_i(t)$$

where $W_1(t), \ldots, W_n(t)$ are standard Brownian motions with correlation $dW_i dW_j = \rho_{ij} dt$. We wish to price a European option, whose terminal payoff is given by $\Phi(S(T))$. Then, we know that the value of the option, $F(t, S(t))$, must satisfy the Black-Scholes PDE:

$$\begin{aligned} & \frac{\partial F}{\partial t} + \sum_{i=1}^{d} r S_i \frac{\partial F}{\partial S_i} + \frac{1}{2} \sum_{i=1}^{d} \sum_{j=1}^{d} S_i S_j \sigma_i \sigma_j \rho_{ij} \frac{\partial^2 F}{\partial S_i \partial S_j} - rF = 0 \\ \text{s.t.} \quad & F(T, S(T)) = \Phi(S(T)) \end{aligned}$$

While our method is capable of solving such a problem for any payoff or dynamics, we will choose a simple example so that we can calculate the solution in closed form as a benchmark. Let us assume that $\mu_i = \mu$, $\sigma_i = \sigma$, and $\rho_{ij} = 0$ for $i \neq j$. We will then price a geometric basket option, whose payoff is given by $\Phi(S(T)) = \max\left\{ \left(\prod_{i=1}^{n} S_i(T)\right)^{\frac{1}{n}} - K, 0 \right\}$. The value of this option has an analytic solution as derived in Vorst (1992).

We fit our model on problems of increasing dimension (n). For each n, we let $r = 5\%$, $\sigma = 20\%$, $T = 5$, and $K = 10$, and our neural-network has four hidden layers of 50 neurons each and uses the tanh activation function.

In this example, we will sample our interior training points from the risk-neural diffusion process, starting with an initial price of $S_0 = (10, 10, \dots)$. The boundary training points are drawn from the distribution of the risk-neutral diffusion process at time T. We could use uniform sampling on a rectangular region, however, as the number of dimensions grows, the probability of the process leaving any bounded box also grows. This leads to declining performance of the model. We again sample 1,000 points for 40,000 iterations for each problem. Because the complexity of the solution is growing in the number of dimensions, it is likely beneficial to increase the complexity of the approximating function or to increase the training time. For comparison purposes, we keep these factors constant as we increase n, and performance is still quite good.

Table 1.1 shows the performance of our method as we increase the number of dimensions. In the first column, we show the percent error at our initial sampling point S_0. The performance of our model declines as we increase the number of dimensions, particularly for 25 or 50 dimensions, but the overall error is still quite small. The second and third columns show the average absolute and percent error for 100,000 points sampled randomly from the same diffusion process as we used in training. The error is again growing, but not exponentially. Finally, the fourth column shows the training time required for each problem running on a single NVIDIA GTX 1070 GPU. The reason that the training time increases is that we must compute more partial derivatives of our approximating function. In particular, the Hessian matrix in the PDE grows on the order n^2. However, we are using uncorrelated processes in our example, so the problem only grows on the order n here. See Sirignano and Spiliopoulos (2017) for a method to approximate the Hessian and speed up training.

We compare our method to the `NDSolve` function in Mathematica. This function uses a finite difference Adams scheme on a grid, so the computation required grows exponentially.[1] We see that for 1 or 2 dimensions, Mathematica is clearly superior. It runs in just a couple of seconds and produces lower error overall. However, for $n = 4$ and greater, the error and computation quickly deteriorate. At $n = 7$, 32GB of ram was insufficient for computation,

[1]We use the default settings. While manual adjustment of method, grid spacing, or approximation order may affect the trade off of accuracy and performance, the pattern is qualitatively similar.

as the 7-dimensional grid was too large to store in memory.

1.3.1.3 Parameters as Dimensions

Another way to increase the dimensionality of the problem is to treat the exogenous parameters r, σ, and K as separate dimensions. Previously, we fixed the values of these parameters before solving for our approximating function. While this approximating function could be reused if the current stock price or time to maturity changes, we would have to fit a new approximation if the other parameters changed. Instead, we will now feed these parameters in to our approximating function as variables.

As before, we will sample $t \in [0, T]$ uniformly. We will also sample $r \in [0, 20\%]$, $\sigma \in [0\%, 30\%]$, and $K \in [0, 20]$. The share price is then sampled from the risk-neutral diffusion corresponding to the chosen parameters. Using the same neural network architecture from the previous section led to poor results, as it seems this function is more complex. Instead, we will use a neural network with five hidden layers of 100 neurons each, which takes longer to train, but can fit more complex functions.

Table 1.3 shows the performance of our approximating function when we fix the parameters at some arbitrary values. The first row is equivalent to the first row of Table 1.1. Despite fitting a more complicated function, the performance on the same slice as before is comparable. Figure 1.3 shows the error on this slice, which is equivalent to Figure 1.2 (d). While this approximate took longer to train, it has much greater reuse value as it can accommodate a wide range of parameter values without needing to retrain.

1.3.1.4 A Method Hazard

In several other papers that use machine learning methods to approximate PDE solutions, the authors generally evaluate the success of the algorithm by the error at a single point, S_0. However, this measure alone can lead to misleading results, as shown below.

Suppose that we followed the same procedure as outlined above, but instead we used the approximating neural network $\hat{f}(t_i; \theta_i)$. Critically, this function *does not* depend on the current stock price. It is only a function of time.

Now if we sample the boundary points from the risk neutral diffusion starting at S_0, our goal for the boundary portion of the loss is to minimize

$$L^{bnd}(\theta_i, A_i^{bnd}) = \sum_{(t_i,S_i)\in A_i^{bnd}} \left(\hat{f}(T;\theta_i) - \Phi(S_T)\right)^2$$

Equivalently, our goal is to select a single value $\hat{f}(T;\theta_i)$ for all values of S_T that minimizes the Mean Squared Error, which we know to be the conditional mean. The function is blind to the current value of S_T, so the most recent information was the initial value S_0. Optimally, $\hat{f}(T;\theta_i) = \mathbb{E}_0[\Phi(S_T)]$

Now for the PDE, all partial derivatives with respect to S_i will be zero, so the PDE is simply $\frac{\partial \hat{f}}{\partial t} - r\hat{f} = 0$. The solution to this given our boundary condition is $\hat{f}(t) = e^{-r(T-t)}\,\mathbb{E}_0[\Phi(S_T)]$, where we have sampled S_T from the risk-neutral diffusion process. When $t = 0$, this corresponds exactly to the result of the Feynman-Kac Theorem.

So, even if our neural network is completely independent of the share price, we will find that it does an excellent job at approximating the true value at the initial point. In applying this method, we must be very careful to examine that the approximation is good over the entire space, not just at a single point.

The results using only $\hat{f}$ are shown in Table 1.4. In the second column, we see that this naive network actually outperforms our prior network from Table 1.1 for many dimensions. However, the average error over the space is quite substantial as shown in columns three and four. Also, the error is decreasing in n, but this is due to the fact that the product of GBMs has decreasing variance as more are added. This leads to less variation in the payoff and a smaller range of prices.

1.3.2 Dynamic Programming and Hamilton-Jacobi-Bellman Equations

In dynamic programming problems, agents are concerned with how to optimally control some processes to minimize the total cost. In finance, these problems often arise as the optimal control of agents seeking to manage a portfolio of assets as in Merton (1973) or manage a firm as in DeMarzo and Sannikov (2006). Analytic solutions to these problems are rare, and only for problems with few dimensions. While finite difference schemes work well in low dimensions, they are limited for problems with a large number of dimensions.

We will consider here a simple Linear-quadratic-Gaussian example for which we can compute the solution relatively simply with the use of Monte Carlo methods. This example is used in Han, Jentzen, and E (2017) as a benchmark for numerical performance of another numerical approximation method.

Suppose we have n state variables, whose dynamics are given by $d\vec{X}_t = 2\sqrt{\gamma}\,\vec{m}_t + \sqrt{2}d\vec{W}_t$ where $\vec{W}_t \in \mathbb{R}^n$ are independent Brownian motions. Let $\vec{X}_0 = [0, \ldots, 0]^\top$, $t \in [0, T]$ and $\vec{X}_t \in \mathbb{R}^n$. $\gamma > 0$ denotes the strength of the control, and $\vec{m}_t \in \mathbb{R}^n$ represents the control process. The associated cost for a given control is given by

$$J(t, \vec{X}_t) = \max_{\vec{m}} \mathbb{E}_t \left[\int_t^T ||\vec{m}_s||_2^2 ds + g(\vec{X}_T) \right]$$

In differential form,

$$0 = \max_{\vec{m}_t} ||\vec{m}_t||_2^2 + \frac{\partial J}{\partial t} + 2\sqrt{\gamma} \sum_{i=1}^{n} \frac{\partial J}{\partial X_{i,t}} m_{i,t} + \sum_{i=1}^{n} \frac{\partial^2 J}{\partial X_{i,t}^2}$$

From the FOC, the optimal control is given by $m_{i,t} = -\sqrt{\gamma}\,\frac{\partial J}{\partial X_{i,t}}$. Plugging this in yields the non-linear PDE that the value process must satisfy:

$$\begin{aligned} 0 &= \frac{\partial J}{\partial t} - \gamma \sum_{i=1}^{n} \left(\frac{\partial J}{\partial X_{i,t}} \right)^2 + \sum_{i=1}^{n} \frac{\partial^2 J}{\partial X_{i,t}^2} \\ \text{s.t.} \quad & J(T, \vec{X}_T) = g(\vec{X}_T) \end{aligned}$$

The solution to this PDE is equal to

$$J(t, \vec{X}_t) = -\frac{1}{\gamma} \ln \left(\mathbb{E}_t \left[\exp \left(-\gamma g(\vec{X}_t + \sqrt{2}\,\vec{W}_T) \right) \right] \right)$$

which can be computed relatively easily using Monte Carlo. While this is true for any function $g(\cdot)$, in this example $g(\vec{x}) = 1 + \ln((1 + ||\vec{x}||_2^2)/2)$ and $\gamma = 1$.

1.3.2.1 One Dimension

As a numerical example, we let $\lambda = 1$ and make 20,000 iterations, sampling 1,000 points on the interior and boundary each time. The learning rate starts at 0.01 and decays 10% every 250 iterations. The approximating function is a standard DNN with four hidden layers of 50 neurons each.

As in the previous example, how to sample from D is an issue. Sampling uniformly from a rectangle again poses problems for values near the spatial boundary. This is illustrated in Figure 1.4 (a). The error is largest for large values of X_0, where the process is again likely to leave the sampled region. However, this problem is much less severe than in the Black-Scholes case, as agents are actively controlling the drift to push the process back to 0. It is only through bad shocks that they are pushed outside. The right figure shows the subset where it would take a three standard deviation shock to push agents past the sampled boundary. Note that the noise in the figure is due to the use of Monte-Carlo simulation as the benchmark. The neural network is very smooth.

Figure 1.4 (b) shows the fit when we instead sample from the uncontrolled Brownian motion (drift of 0) starting at $X_0 = 0$. The quality of the fit does not improve much over the subsample region. Despite not wasting points on undetermined regions, it doesn't appear to offer much gain. However, the problem of leaving the boundary set grows as the dimensions increase, so there are gains in higher dimensions.

Finally, Figure 1.4 (c) shows the fit when the samples are drawn from the controlled Brownian motion, under the current estimate of the optimal control. This understandably leads to poor fit for extreme values of X_T, as the process has learned to avoid those regions. Hence, few points are sampled there to discipline the approximation. While this method does appear to yield good solutions at $X_0 = 0$, this may not always be the case. When the boundary condition $g(x)$ is non-monotonic, it may be that the initial estimate samples may points from a local minimum, and not the global minimum. It will continue to sample points along this path and may never discover the global minimum, leading to erroneous values.

When sampling from the controlled Brownian motion, this problem effectively becomes a reinforcement learning problem. There are vasts amounts of research in this field, where agents learn by doing and act according to their current model of the value function (Sutton and Barto (2018)). This is also the method adopted by Han, Jentzen, and E (2017) for PDE solutions, but only produces solutions at a single starting point.

1.3.2.2 Multiple Dimensions

We will now extend this example to problems of increasing dimensionality. As before, let $\lambda = 1$ and make 20,000 iterations, sampling 1,000 points from the uncontrolled Brownian motion in the interior and boundary each time. The learning rate starts at 0.01 and decays 10% every 250 iterations. The approximating function is a standard DNN with four hidden layers of 50 neurons each.

Table 1.5 shows the performance of this method. The percent error is fairly constant as the dimensionality increases, while the computation time grows linearly. This is in comparison to the performance of Mathematica's finite difference method, seen in Table 1.6, where the error grows considerably beyond three dimensions. Again, 32GB of memory was insufficient for seven dimensions in Mathematica.

1.3.3 Mean Field Games

1.3.3.1 Traffic Problem: 1-Dimension

Here, we consider a simple example model where agents seek to avoid traffic during their commute. This problem is taken from Alanko (2015), who uses it to test an alternative numerical method. While the dimensionality of this problem is low, it is difficult to solve using traditional techniques due to non-linearities in the PDE. Additionally, the boundary conditions take the form of an initial condition for the density function, but a terminal condition for the value function.

Suppose pedestrians are distributed along a line, with initial distribution $m_0(x) = \mathcal{N}(-1, 0.25^2)$. That is, at time 0, they are normally distributed around a town at $x = -1$. By the end of the game at time T, they would like to be located at their home town at $x = 1$. Pedestrians control their motion as

$$dX_t = \alpha(X_t, t)dt + dW_t$$

where α is their chosen travel velocity. Their goal is to minimize the expected cost

$$u(X_t, t) = \mathbb{E}\left[g(X_T) + \int_t^T \frac{\alpha_s^2}{2} ds + \int_t^T m(X_s, s)ds\right]$$

where the terminal cost is given by $g(x) = e^{-2x^2} - 2e^{-2(x-1)^2}$.

The terminal cost penalizes pedestrians caught between the two towns at the end, and rewards those who make it home to $x = 1$. The expected cost also contains the cost of movement choice, which is $\frac{\alpha_s^2}{2}$. Finally, the cost also contains a penalty for being at the same position on the road as many other pedestrians, $m(X_s, s)$.

After substituting in the optimal control, we obtain the system of PDEs:

$$\frac{\partial u}{\partial t} - \frac{1}{2}\left(\frac{\partial u}{\partial x}\right)^2 + m(x,t) + \frac{1}{2}\frac{\partial^2 u}{\partial x^2} = 0, \qquad u(x,T) = g(x) \tag{1.6}$$

$$\frac{\partial m}{\partial t} - \frac{\partial}{\partial x}\left(m(x,t)\frac{\partial u}{\partial x}\right) - \frac{1}{2}\frac{\partial^2 m}{\partial x^2} = 0, \qquad m(x,0) = m_0(x) \tag{1.7}$$

To solve this system, we will model u as a standard DNN. However, $m(x,t)$ is a distribution function, so we will model it as a mixture density network which we described in Section 1.2.2.5. The details are described below:

0. $i = 0$.

    - Initialize our approximating cost function $\hat{u}$ as standard DNN with three hidden layers of 50 neurons each: $\hat{u}(x, t; \theta_0^u)$ with random parameters θ_0^u.
    - Initialize our approximating distribution function $\hat{m}$ as a mixture density network with one hidden layer of 50 neurons feeding into 5 normal distributions: $\hat{m}(x, t; \theta_0^m)$ with random parameters θ_0^m.

1. Draw sample of interior points $\hat{D}_i$. Sample $(t_j, S_j) \in [0,T] \times [\underline{x}, \overline{x}]$ uniformly.

2. Compute the PDE residuals at these points for both PDEs.

$$R_1^u(t_j, x_j, \hat{u}(t_j, x_j; \theta_i^u)) = \frac{\partial \hat{u}}{\partial t} - \frac{1}{2}\left(\frac{\partial \hat{u}}{\partial x}\right)^2 + \hat{m} + \frac{1}{2}\frac{\partial^2 \hat{u}}{\partial x^2}$$
$$R_1^m(t_j, x_j, \hat{m}(t_j, x_j; \theta_i^m)) = \frac{\partial \hat{m}}{\partial t} - \frac{\partial}{\partial x}\left(\hat{m}\frac{\partial \hat{u}}{\partial x}\right) - \frac{1}{2}\frac{\partial^2 \hat{m}}{\partial x^2}$$

3. • Draw sample of boundary points $\hat{B}_i^u$. Sample $(t_j, x_j) \in [T, T] \times [\underline{x}, \overline{x}]$ uniformly.
   • Draw sample of boundary points $\hat{B}_i^m$. Sample $(t_j, x_j) \in [0, 0] \times [\underline{x}, \overline{x}]$ uniformly.

4. Compute the boundary residuals for both boundaries on their respective points.

$$R_2^u(t_j, x_j, \hat{u}(t_j, x_j; \theta_i^u)) = \hat{u}(t_j, x_j; \theta_i^u) - g(x_j)$$
$$R_2^m(t_j, x_j, \hat{m}(t_j, x_j; \theta_i^m)) = \hat{m}(t_j, x_j; \theta_i^m) - m_0(x_j)$$

5. Compute the total loss.

$$\begin{aligned} L(\theta_i^u, \theta_i^m; \hat{D}_i, \hat{B}_i^u, \hat{B}_i^m) = & \sum_{(t_j,x_j)\in\hat{D}_i} \left(R_1^u(t_j, x_j, \hat{u}(t_j, x_j; \theta_i^u))^2 + R_1^m(t_j, x_j, \hat{m}(t_j, x_j; \theta_i^m))^2\right) \\ & + \lambda \sum_{(t_j,x_j)\in\hat{B}_i^u} R_2^u(t_j, x_j, \hat{u}(t_j, x_j; \theta_i^u))^2 \\ & + \lambda \sum_{(t_j,x_j)\in\hat{B}_i^m} R_2^m(t_j, x_j, \hat{m}(t_j, x_j; \theta_i^m))^2 \end{aligned}$$

6. If the total loss is sufficiently small, end. Otherwise, update the parameters of our approximating function (assign $\theta_{i+1}^u, \theta_{i+1}^m$) using the ADAM optimizer to minimize the total loss. Update the learning rate, set $i = i + 1$, and return to step 1.

For this example, $\lambda = 1$, $T = 1$, and x is sampled from $[\underline{x}, \overline{x}] = [-6, 6]$. The model was trained for 10,000 iterations with each sampling 1,000 points per set. The initial learning rate was 0.005, decaying 10% every 500 iterations.

The results are shown in Figure 1.5. We see that the boundary conditions are well satisfied, and that we get the expected behavior of pedestrians moving to the home town. A finite difference reference scheme is unfinished as of yet, but our results are highly consistent with the finite difference solution shown in Alanko (2015).

1.3.3.2 Traffic Problem: 2-Dimensions

The above problem can be expanded to two spatial dimensions as follows (again taken from Alanko (2015)). Let $\vec{x} = [x_1, x_2]$ and locate the starting town and home town at $(x_1, x_2) = (-1, 0)$ and $(x_1, x_2) = (1, 0)$ respectively. The pedestrians initial distribution is given by $m_0(\vec{x}) = \mathcal{N}((-1, 0), 0.25^5 \mathbb{I}_2)$. Pedestrians movement is governed by

$$d\vec{X}_t = \vec{\alpha}(\vec{X}_t, t)dt + d\vec{W}_t$$

where $\vec{\alpha} \in \mathbb{R}^2$ is again the chosen velocity, now in two dimensions. The agents seek to minimize their expected cost

$$u(\vec{X}_t, t) = \mathbb{E}\left[g(\vec{X}_T) + \int_t^T \frac{||\vec{\alpha}_s||_2^2}{2} ds + \int_t^T m(\vec{X}_s, s) ds + \int_t^T X_{2,s}^2 ds\right]$$

where the terminal cost is given by $g(\vec{x}) = e^{-2(x_1^2 + x_2^2)} - 2e^{-2(x_1-1)^2 + x_2^2}$.

The terminal cost is the two dimensional equivalent of the previous cost, which penalizes pedestrians caught between the two towns at the end, and rewards those who make it home. The expected cost also contains the cost of movement cost in both directions and a penalty for being at the same position on the road as many other pedestrians. Additionally, the last term penalizes pedestrians from straying off the road, which is the x_1 axis.

After substituting in the optimal control, we obtain the system of PDEs:

$$\frac{\partial u}{\partial t} - \frac{1}{2}\left[\left(\frac{\partial u}{\partial x_1}\right)^2 + \left(\frac{\partial u}{\partial x_2}\right)^2\right] + m(\vec{x}, t) + x_2^2 + \frac{1}{2}\left(\frac{\partial^2 u}{\partial x_1^2} + \frac{\partial^2 u}{\partial x_2^2}\right) = 0, \qquad u(\vec{x}, T) = g(\vec{x})$$

$$\frac{\partial m}{\partial t} - \frac{\partial}{\partial x_1}\left(m(\vec{x}, t)\frac{\partial u}{\partial x_1}\right) - \frac{\partial}{\partial x_2}\left(m(\vec{x}, t)\frac{\partial u}{\partial x_2}\right) - \frac{1}{2}\left(\frac{\partial^2 m}{\partial x_1^2} + \frac{\partial^2 m}{\partial x_2^2}\right) = 0, \quad m(\vec{x}, 0) = m_0(\vec{x})$$

The solution method for this problem is analogous to the 1-dimensional case. For this example, $\lambda = 1$, $T = 1$, and $\vec{x}$ is sampled from $[\underline{x}, \overline{x}]^2 = [-6, 6]^2$. The model was trained for 100,000 iterations with each sampling 1,000 points per set. The initial learning rate was 0.01, decaying 2% every 500 iterations.

The results are shown in Figure 1.6. The finite difference verification in a work in progress, but the results are highly consistent with the finite difference solution in Alanko (2015).

1.4 Conclusion

Our paper has outlined a numerical method for solving high dimensional PDEs accurately and efficiently. By using deep neural networks as approximating functions, it is possible to take advantage of highly efficient machine learning techniques. The backpropagation algorithm allows for rapid computation of derivatives for use in the PDE itself and of gradients for updating parameters, even when the problem is non-linear. Batch gradient descent substitutes costly computation of the PDE over the entire domain with gradient approximations on a sampled subset, removing the need for memory intensive point grids. Additionally, restrictions that functions must satisfy the properties of a probability density function can be easily accommodated with the use of a mixture density network as the approximating function.

Our method was highly accurate for the Black-Scholes PDE in up to 50 spatial dimensions. For Black-Scholes pricing of basket options, our method had error below 0.3% on average over a wide range of possible prices with up to 15 assets. The accuracy deteriorated somewhat with 25 or 50 dimensions, but was still acceptable and could be improved with a larger DNN. We were also able to accurately approximate the solution when the exogenous parameters were free inputs, which provides greater reusability. However, it is important to consider that even an approximation blind to the stock price can still produce good estimates at a single point, as this PDE devolves to the Feynman-Kac method.

Similarly, our method was highly accurate for non-linear dynamic programming PDEs in high dimensions. The average error over the space remained constant, below 1%, as the number of dimensions grew. The finite difference comparison struggled beyond three dimensions. However, we illustrated the potential shortcoming of relying on approximations beyond the region pinned down by the boundary conditions. Unlike simple finite different schemes that may be undefined on problem regions, the neural network will return inaccurate estimates.

Our method also successfully modeled the solution to one and two dimensional mean field game problems which incorporate a complex system of constraints. By using a mixture density network as the approximating function, we were able to automatically satisfy the restriction that one solution was a valid PDF.

Together, this method can allow finance researchers to explore richer models than previous techniques. This is particularly useful for mean field games, models featuring many economic state variables, or models with many agents who are heterogeneous or have non-aggregating preferences.

Finally, because this solution method is built on top of popular machine learning frameworks, we are able to leverage the easy to use and highly optimized software libraries like TensorFlow from Google. This library handles all of the derivative computation and optimizer updating in the background, allowing researchers to quickly iterate through many model variations with minimal coding.

Figure 1.1: Fixed sampling may ignore problem points

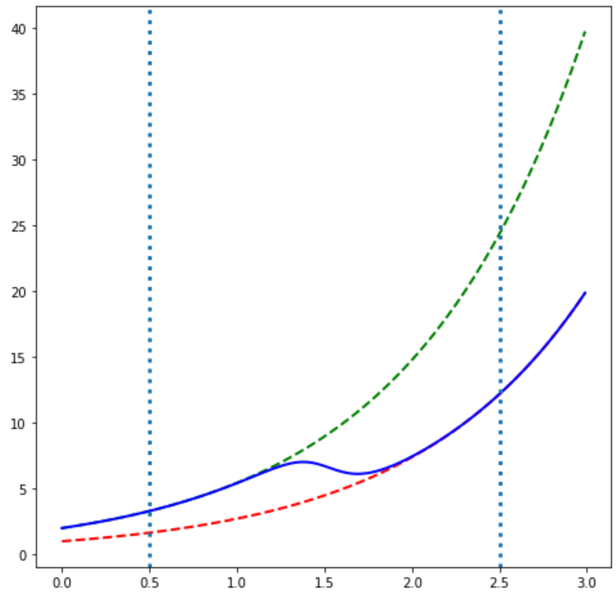

We see two possible solutions to the ODE $f'(x) = f(x)$ with different initial boundary conditions (dashed lines). The solution we would like is the dashed green line. However, we only sampled two points at the vertical blue lines, $x = 0.5$ and $x = 2.5$. The PDE residuals at the sampled points are very small, and the boundary condition is satisfied. This problem is only resolved by sampling points around $x = 1.5$ to correct the error.

Dimensions (n)	Percent Error at S_0	Average Absolute Error	Average Percent Error	Time (Seconds)
$n = 1$	0.002%	0.0015	0.183%	107
$n = 2$	0.008%	0.0011	0.159%	134
$n = 3$	0.007%	0.0020	0.262%	159
$n = 4$	0.110%	0.0034	0.465%	170
$n = 5$	0.018%	0.0012	0.174%	203
$n = 10$	0.013%	0.0021	0.265%	331
$n = 15$	0.028%	0.0022	0.250%	456
$n = 25$	0.233%	0.0106	1.087%	660
$n = 50$	0.294%	0.0423	3.385%	1251

Table 1.1: Performance of Neural Network on Multi-dimensional Black-Scholes

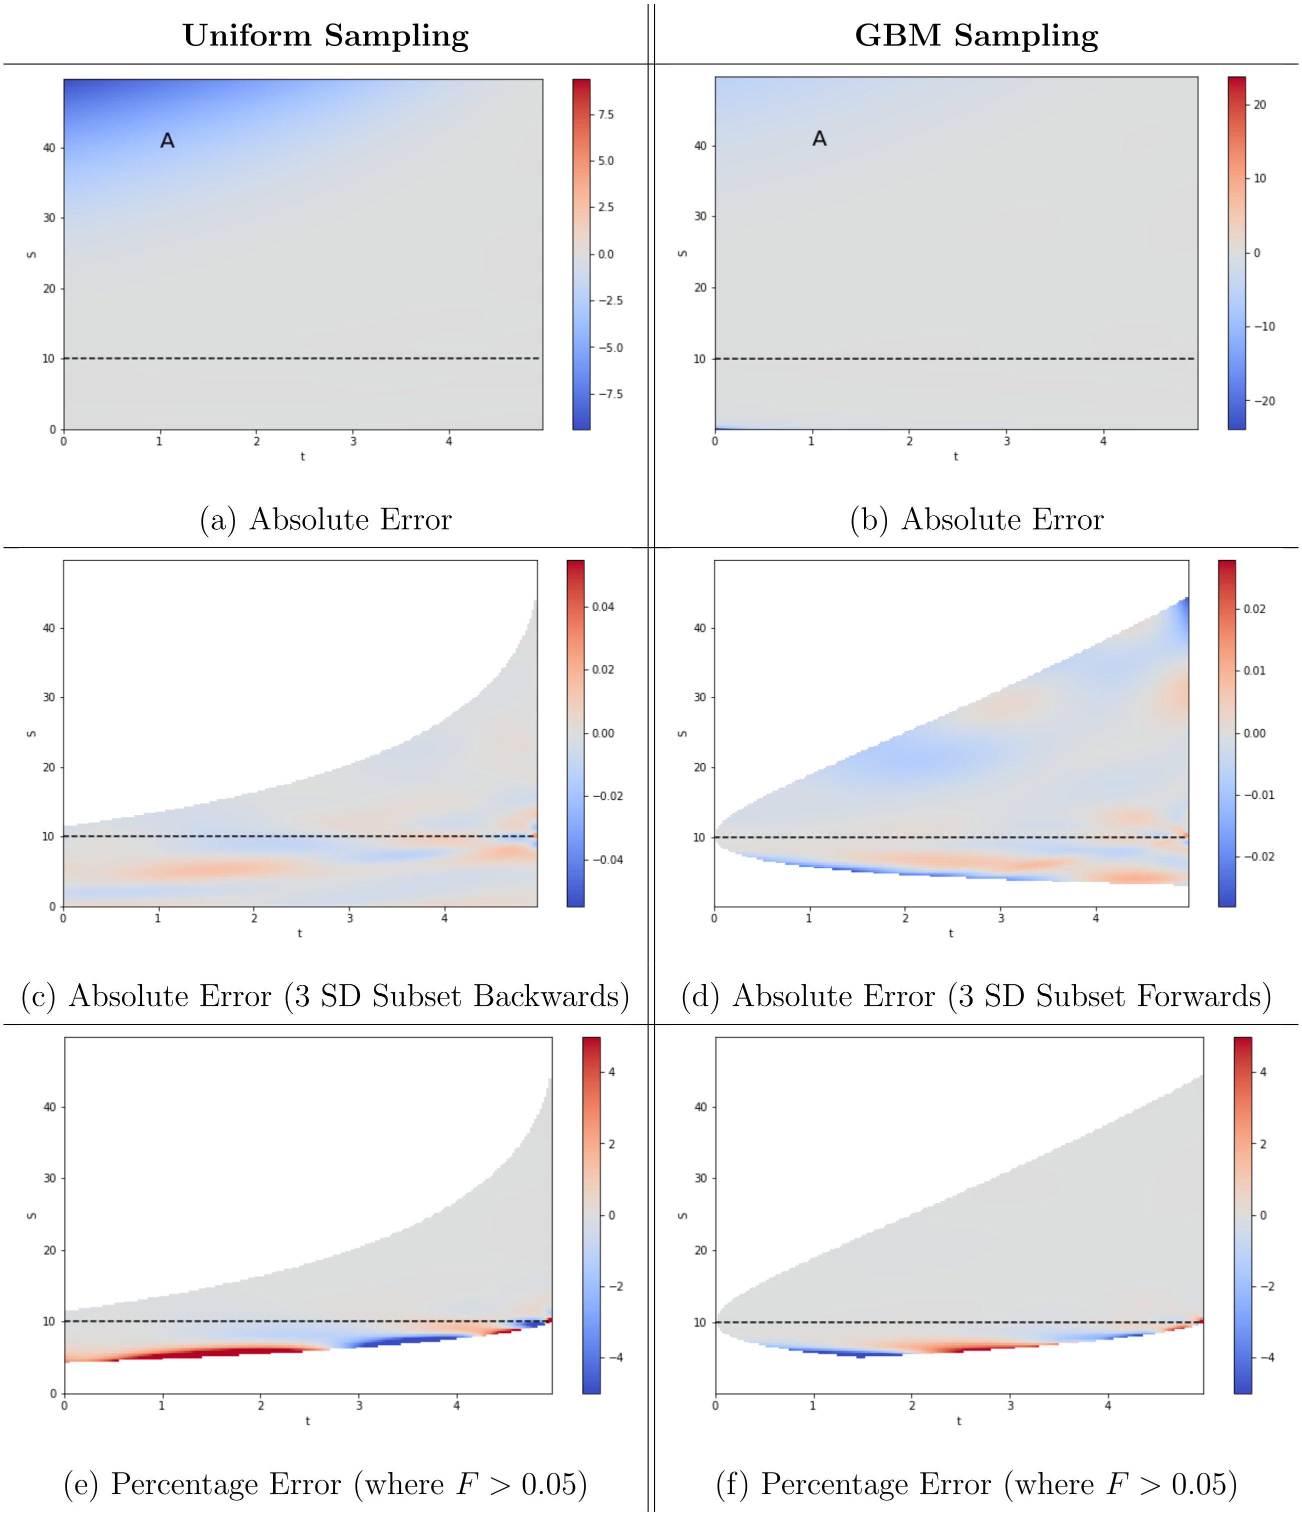

Uniform Sampling | **GBM Sampling**

(a) Absolute Error

(b) Absolute Error

(c) Absolute Error (3 SD Subset Backwards)

(d) Absolute Error (3 SD Subset Forwards)

(e) Percentage Error (where $F > 0.05$)

(f) Percentage Error (where $F > 0.05$)

Figure 1.2: Error for the European Call Option

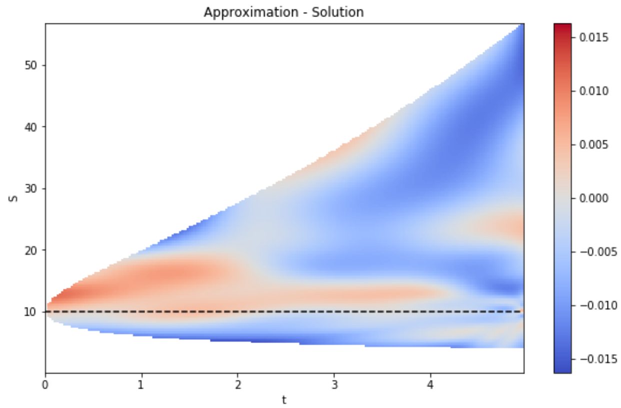

Figure 1.3: Absolute Error of Neural Network on Black-Scholes with Variable Parameters ($r = 10\%, \sigma = 20\%, K = 10$)

Dimensions (n)	Percent Error at S_0	Average Absolute Error	Average Percent Error	Time (Seconds)
$n = 1$	0.002%	0.0002	0.010%	1
$n = 2$	0.003%	0.0000	0.004%	11
$n = 3$	0.035%	0.0028	0.416%	33
$n = 4$	1.295%	0.0248	3.343%	51
$n = 5$	1.622%	0.0244	2.623%	499
$n = 6$	1.445%	0.0223	2.319%	5,389
$n = 7$	Failed			

Table 1.2: Performance of Mathematica `NDSolve` on Multi-dimensional Black-Scholes

r	σ	K	Percent Error at S_0	Average Absolute Error	Average Percent Error
10%	20%	10	0.076%	0.0035	0.209%
5%	25%	5	0.122%	0.0071	0.138%
0%	30%	0	0.213%	0.0158	0.154%
15%	15%	10	-0.094%	0.0068	0.347%

Table 1.3: Performance of Neural Network on Black-Scholes with Variable Parameters

Absolute Error **Absolute Error (3 SD Subset Backwards)**

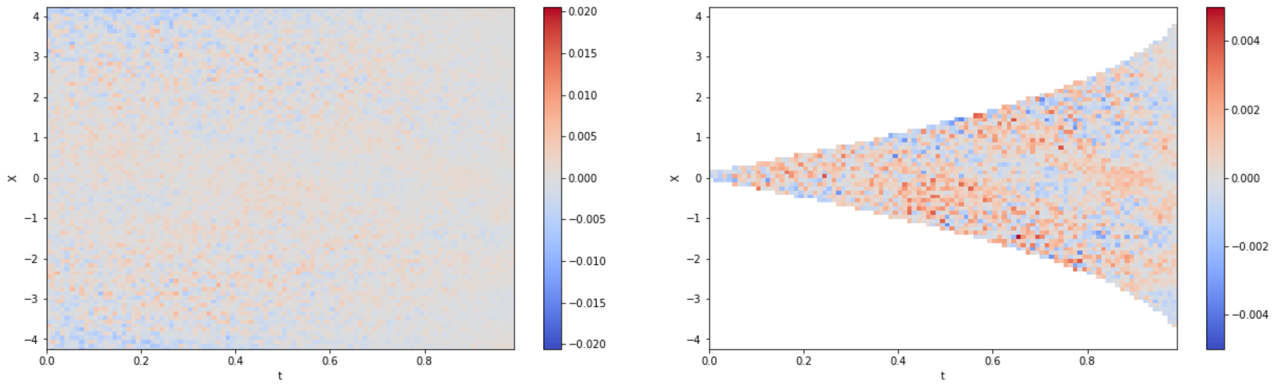

(a) Uniform Rectangular Sampling

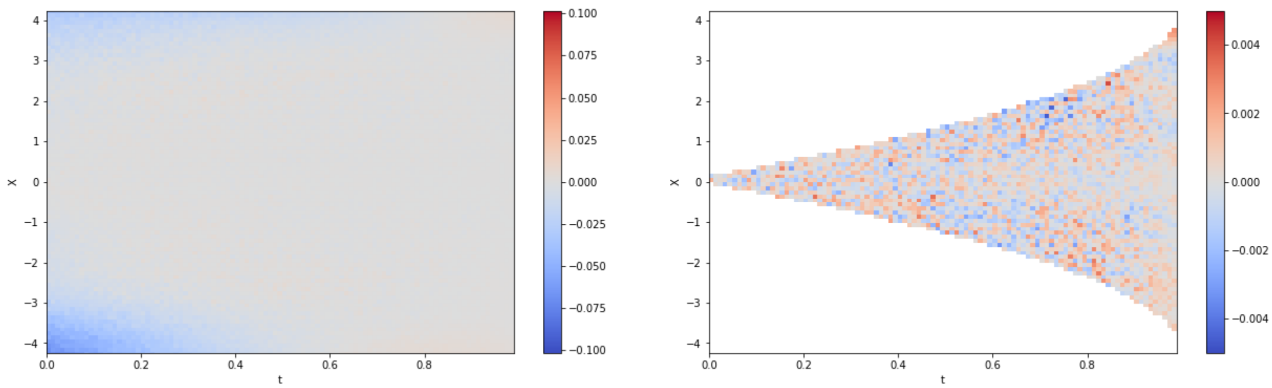

(b) Uncontrolled BM Sampling

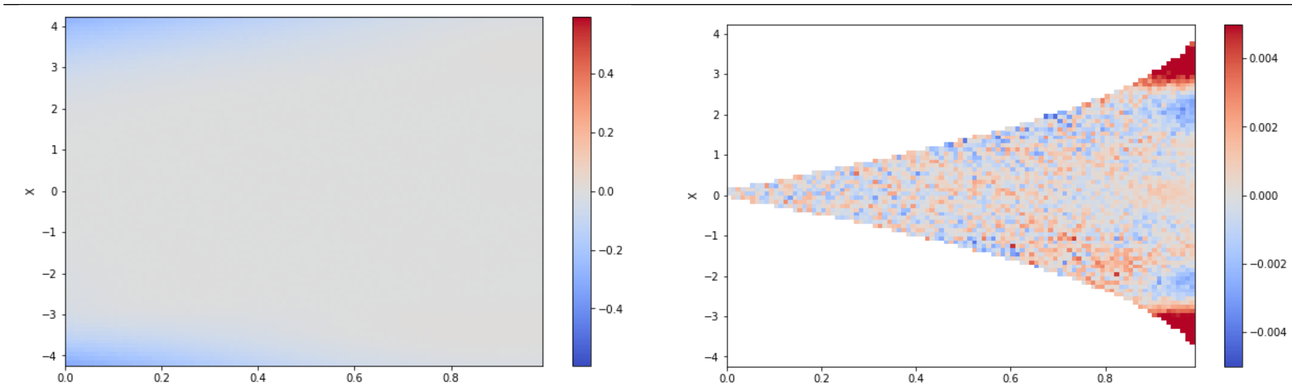

(c) Controlled BM Sampling

Figure 1.4: Error for 1-Dimensional HJB Equation (Noise from Monte-Carlo Reference)

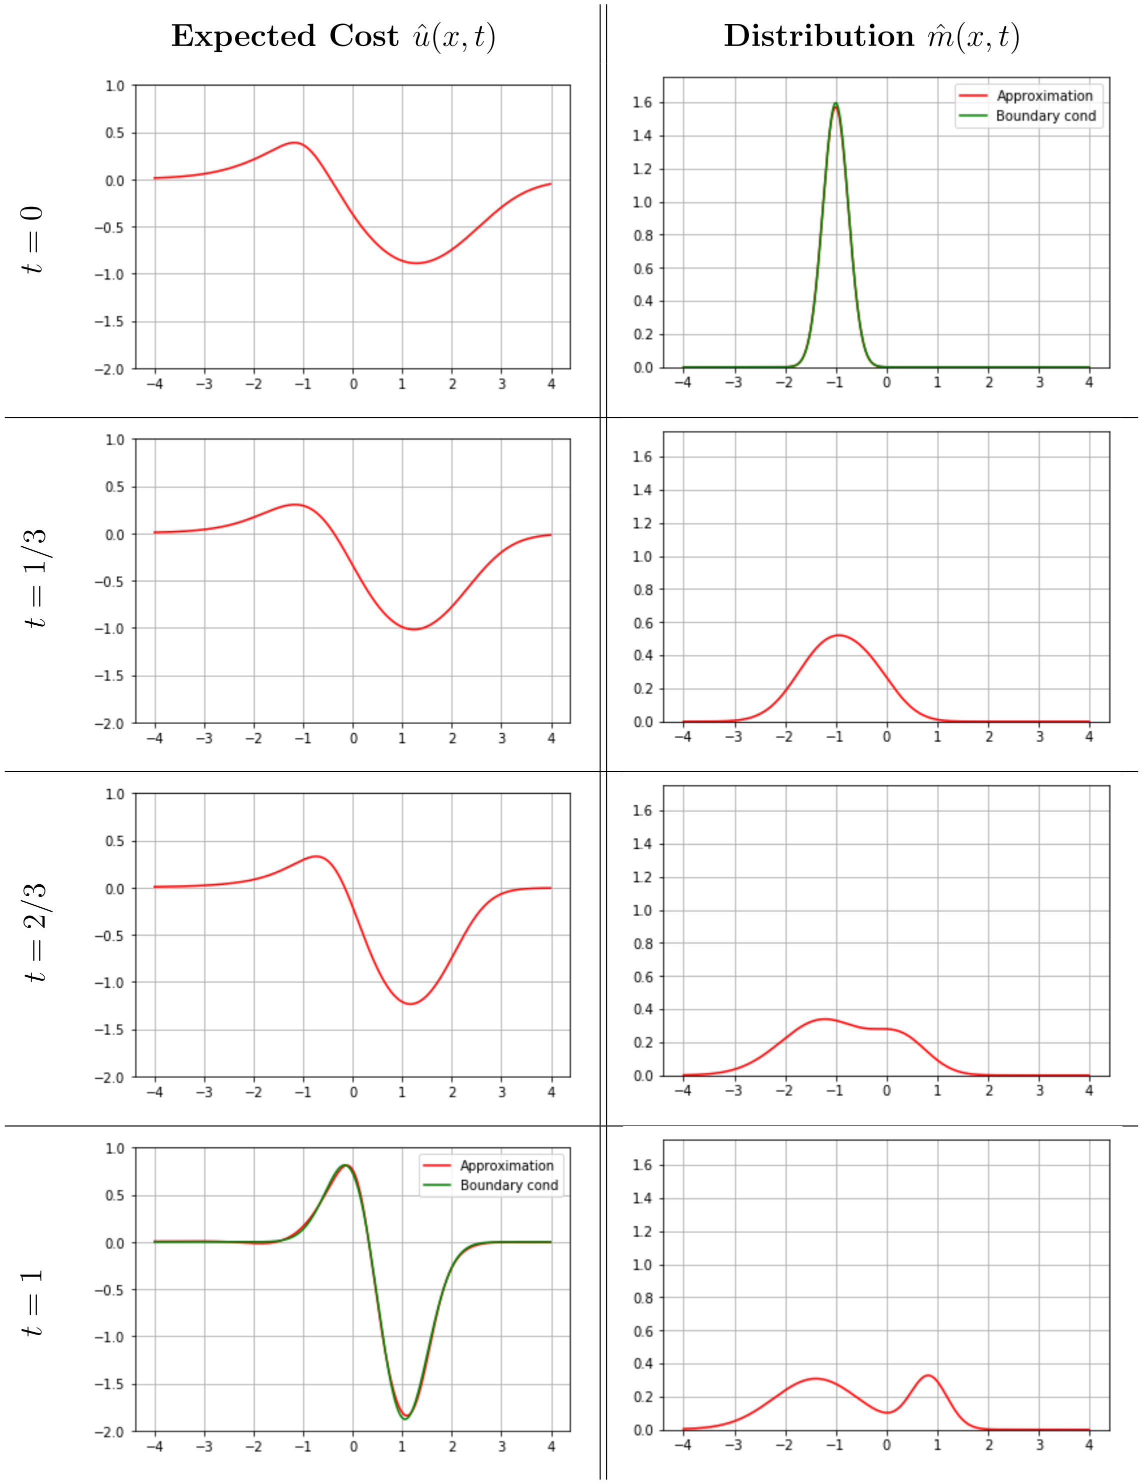

Figure 1.5: 1-Dimensional Mean Field Game

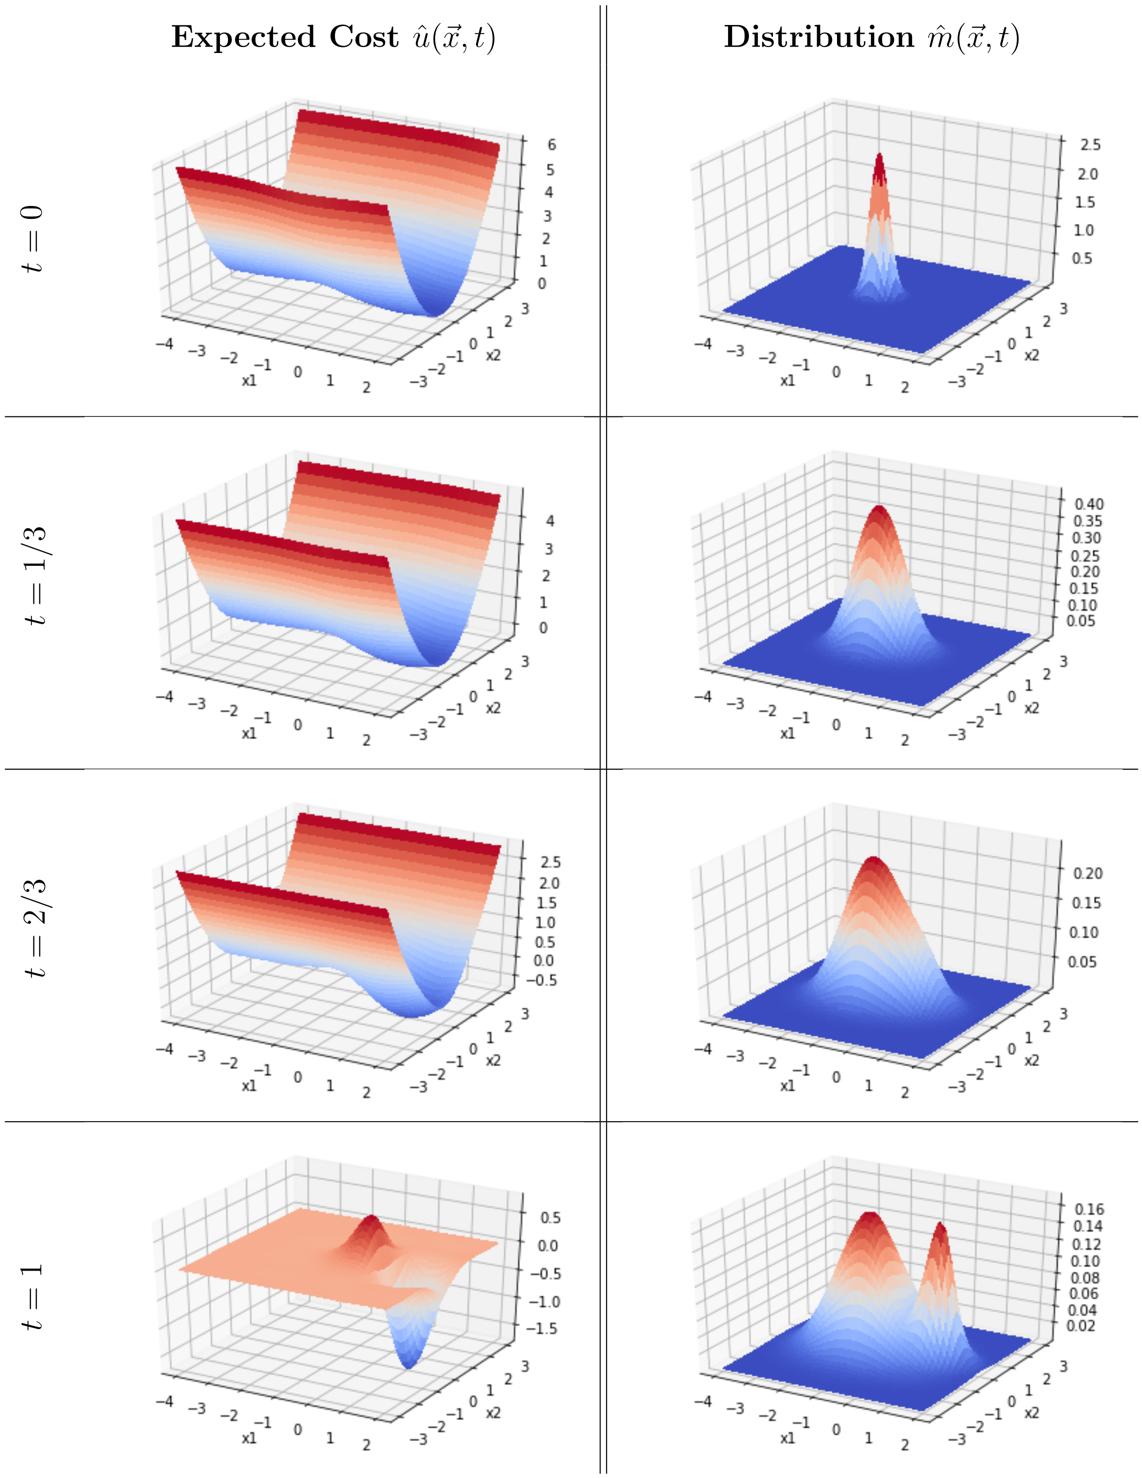

Figure 1.6: 2-Dimensional Mean Field Game

Dimensions (n)	Percent Error at S_0	Average Absolute Error	Average Percent Error
$n = 1$	0.170%	2.2479	256.887%
$n = 2$	0.127%	1.5254	197.444%
$n = 3$	0.017%	1.2544	168.696%
$n = 4$	0.062%	1.0884	149.895%
$n = 5$	0.058%	0.9897	138.539%
$n = 10$	0.001%	0.7329	100.784%
$n = 15$	0.692%	0.6121	80.106%
$n = 25$	0.258%	0.4913	58.564%
$n = 50$	0.005%	0.3517	33.671%

Table 1.4: Performance of Neural Network Using Only t on Multi-dimensional Black-Scholes

Dimensions (n)	Percent Error at $\vec{X}_0$	Average Absolute Error	Average Percent Error	Time (Seconds)
$n = 1$	0.071%	0.0155	1.236%	46
$n = 2$	0.126%	0.0179	0.908%	60
$n = 3$	0.515%	0.0185	0.738%	73
$n = 4$	0.274%	0.0180	0.619%	77
$n = 5$	0.006%	0.0177	0.560%	89
$n = 10$	0.086%	0.0306	0.758%	145
$n = 15$	0.069%	0.0391	0.881%	201
$n = 25$	0.040%	0.0212	0.432%	312
$n = 50$	0.075%	0.0381	0.679%	593

Table 1.5: Performance of Neural Network on Multi-dimensional HJB

Dimensions (n)	Percent Error at $\vec{X}_0$	Average Absolute Error	Average Percent Error	Time (Seconds)
$n = 1$	0.751%	0.0038	0.411%	1
$n = 2$	0.129%	0.0046	0.332%	2
$n = 3$	0.156%	0.0052	0.298%	7
$n = 4$	5.499%	0.1118	6.001%	37
$n = 5$	3.572%	0.0853	4.028%	215
$n = 6$	2.080%	0.0662	2.845%	1605
$n = 7$	Failed			

Table 1.6: Performance of Mathematica `NDSolve` on Multi-dimensional HJB

1.5 Bibliography

Achdou, Yves, Francisco J Buera, et al. (2014). "Partial differential equation models in macroeconomics." *Phil. Trans. R. Soc. A* 372.2028, p. 20130397.

Achdou, Yves, Jiequn Han, et al. (2017). *Income and wealth distribution in macroeconomics: A continuous-time approach.* Tech. rep. National Bureau of Economic Research.

Alanko, Samu (2015). "Regression-based Monte Carlo methods for solving nonlinear PDEs." *ProQuest Dissertations Publishing.*

Bishop, Christopher M. (1994). "Mixture Density Networks." *The effects of brief mindfulness intervention on acute pain experience: An examination of individual difference.*

Black, Fischer and Myron Scholes (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81.3, pp. 637–654.

Boyd, John P (2001). *Chebyshev and Fourier spectral methods.* DOVER Publications, Inc.

Brumm, Johannes and Simon Scheidegger (2017). "Using adaptive sparse grids to solve high-dimensional dynamic models." *Econometrica* 85.5, pp. 1575–1612.

Brunnermeier, Markus K and Yuliy Sannikov (2016). "Macro, money, and finance: A continuous-time approach." *Handbook of Macroeconomics.* Vol. 2. Elsevier, pp. 1497–1545.

Cox, John C et al. (1985). "A Theory f the Term Structure of Interest Rates." *Econometrica.*

Cybenko, G (1989). "Approximation by superpositions of a sigmoidal function." *Mathematics of Control, Signals and Systems* 2.4, pp. 303–314.

DeMarzo, Peter M and Yuliy Sannikov (2006). "Optimal security design and dynamic capital structure in a continuous-time agency model." *The Journal of Finance* 61.6, pp. 2681–2724.

Dissanayake, M W and N Phan-Thien (1994). "Neural-network-based approximations for solving partial differential equations." *Communications in Numerical Methods in Engineering* 10.3, pp. 195–201.

Duarte, Victor (2018). "Machine Learning for Continuous-Time Economics." PhD thesis. MIT.

Fernandez-Villaverde, Jesus, Juan Francisco Rubio-Ramirez, and Frank Schorfheide (2016). "Solution and estimation methods for DSGE models." *Handbook of Macroeconomics*. Vol. 2. Elsevier, pp. 527–724.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. MIT Press.

Han, Jiequn, Arnulf Jentzen, and Weinan E (2017). "Solving high-dimensional partial differential equations using deep learning." *Proceedings of the National Academy of Sciences of the United States of America* 115.34, pp. 8505–8510.

Hornik, Kurt (1991). "Approximation capabilities of multilayer feedforward networks." *Neural Networks* 4.2, pp. 251–257.

Judd, Kenneth L (1998). *Numerical methods in economics*. MIT press.

Kingma, Diederik P. and Jimmy Ba (2014). "Adam: A Method for Stochastic Optimization."

Kumar, Manoj and Neha Yadav (2011). "Multilayer perceptrons and radial basis function neural network methods for the solution of differential equations: A survey." *Computers and Mathematics with Applications*.

Lagaris, Isaac Elias, Aristidis Likas, and Dimitrios I. Fotiadis (1998). "Artificial neural networks for solving ordinary and partial differential equations." *IEEE Transactions on Neural Networks*.

Lasry, Jean-Michel and Pierre-Louis Lions (2007). "Mean field games." *Japanese journal of mathematics* 2.1, pp. 229–260.

Leland, Hayne E (1994). "Corporate debt value, bond covenants, and optimal capital structure." *The journal of finance* 49.4, pp. 1213–1252.

Merton, Robert C (1973). "An intertemporal capital asset pricing model." *Econometrica: Journal of the Econometric Society*, pp. 867–887.

Robbins, Herbert and Sutton Monro (1951). "A Stochastic Approximation Method." *The Annals of Mathematical Statistics* 22.3, pp. 400–407.

Sannikov, Yuliy (2008). "A continuous-time version of the principal-agent problem." *Review of Economic Studies*.

Sirignano, Justin and Konstantinos Spiliopoulos (2017). "DGM: A deep learning algorithm for solving partial differential equations."

Snoek, Jasper, Hugo Larochelle, and Ryan P Adams (2012). "Practical bayesian optimization of machine learning algorithms." *Advances in neural information processing systems*, pp. 2951–2959.

Sutton, Richard S and Andrew G Barto (2018). *Reinforcement learning: An introduction.* MIT press.

Vorst, Ton (1992). "Prices and hedge ratios of average exchange rate options." *International Review of Financial Analysis* 1.3, pp. 179–193.

Winschel, Viktor and Markus Kratzig (2010). "Solving, estimating, and selecting nonlinear dynamic models without the curse of dimensionality." *Econometrica* 78.2, pp. 803–821.

You, Seungil et al. (2017). "Deep lattice networks and partial monotonic functions." *Advances in Neural Information Processing Systems*, pp. 2981–2989.

CHAPTER 2

Unenforced Endogenous Reciprocity

with Ivo Welch

2.1 Introduction

Residents of small towns routinely purchase goods from local vendors at prices higher than those at which they could purchase elsewhere. A contractor may purchase her new car from the local dealer, even when the dealer in the next town over is offering a somewhat better deal. In many cases, this can be explained by lower transaction costs, explicit or implicit contracts, better service, or after-market interactions. Reciprocity can be and often is facilitated with strong coordination mechanisms, such as contracts or sanctions (such as tit-for-tat strategies, as in Axelrod (1984).) Yet, it is neither always feasible nor always an accurate description of reality that neighbors can observe one another and/or use contracting mechanisms.

Instead, neighborly reciprocity may be sustained by a weaker mechanism. Buyers can be better off if they voluntarily take into account that their own purchases can enhance economic activity in their neighborhoods. *This is because local surplus is more likely to come back to them in turn.* The contractor may expect the local car dealer to be more likely to use some of his resulting surplus to build another showroom. Some local feedback could also be indirect (and thus more difficult to contract)—the contractor could also benefit when the car dealer sees a local physician, who in turn becomes more inclined to build her new mansion, which in turn could preferentially favor the local contractor.

In our model, agents internalize that they should make decisions not just based on goods and prices, but also based on the return externality. "Purchasing local" is a mechanism to capture the positive "return" externality.[1] We show that such unenforced reciprocity can sustain itself in a small neighborhood within a network, even in the absence of other stronger coordination mechanisms. Agents voluntarily choose to strengthen their local economies by buying local—at least to a point. As in the public goods and as in the group selection literatures, this is plausible only in situations in which the scale of coordination remains small, where the return benefits are large, and (implicitly) where the costs of alternative mechanisms are higher. Voluntary reciprocity can be a very cheap mechanism. When repeat interactions are likely, we find that the reciprocity equilibrium outcome can be at or near first-best.

Despite its intuitive appeal, voluntary reciprocity is a difficult modeling problem. It requires multiple but not a large number of agents. It requires seller surplus, so that sellers are not indifferent to neighborly preferences. It requires the presence of wealth effects, so that agents can expect to become better off when their neighbors are wealthier. It requires a higher probability of capturing surplus from a neighbor than from a non-neighbor, and without last-period selfish unraveling. And it requires buyers that can handicap sellers (neighbors vs. non-neighbors), and sellers that decide on participation and pricing strategies, given these buyers' handicaps.

Individual surplus and reciprocity are mutually necessary in our model. It is only the expected reciprocity that allows sellers to earn a surplus, and it is only the surplus that allows reciprocity to be sustainable. Moreover, the model suggests that unenforced reciprocity can not only be individually but also socially beneficial. The resulting coordination can help reduce randomly duplicated production or randomly missing production.

[1] Externalities are often called "market failures." Here, pure price competition may be a market failure. Similar failures have been observed in the group selection literature (e.g., Wilson and Sober (1994)), where selfish individuals are unable to recapture positive group externalities. However, our mechanism does not require kinship. There is also a tie-in to the theory of the firm. Neither contracts (and enforcement) nor merging may be required to induce parties to coordinate effectively. Firms voluntarily ignore market prices, at least up to a point.

2.2 The Model

Our paper takes a first step to construct a model in which unenforced reciprocal behavior can thrive—where buyers may not "selfishly" defect and buy from the lowest-cost vendor. But when will buyers endogenously act neighborly and local? When would sellers not simply enjoy the higher sales price now and not return the favor when it is their turn (i.e., by buying from the cheaper vendors elsewhere)? If there is reciprocal behavior that induces buyers to favor locals, when would other sellers choose to remain in the market? Are there advantages to reciprocity? What are the limits to reciprocity?

2.2.1 Setup and Assumptions

We assume that sellers must first produce and then offer their goods for sale at a posted price. A good that is not sold declines in value. Such markets are among the most common economic arrangements today: cars are manufactured first and then offered in show rooms. A car that is not sold early must be more aggressively discounted. A contractor who has not procured a job can never get his time back. More specifically, our model assumes that products that are not sold become worthless.

Specifically, we model two potential sellers who can make period-by-period production and pricing decisions. They bring their goods to market only when they expect not to lose money in this period. A buyer can then inspect the sellers' prices and purchase from the cheapest seller *after* applying an (endogenous) handicap. All sellers know the buyer's handicaps before they produce (or not). The handicaps can be identical, in which case the buyer treats the sellers the same. Or, the buyer can handicap local sellers differently. It is this competition among sellers, subject to the public (endogenous) buyer handicaps, that will pin down (endogenous) market participation and product pricing in our model.

The sellers' game has no pure strategy equilibria. If the sellers' participation and prices were deterministic, then each seller who would underbid the other sellers' known (handicapped) prices by epsilon would always win the business. The losing sellers would always have paid the participation cost and yet have come out empty. But then no seller would enter. Thus, the only viable equilibria will be mixed. The sellers can randomize (entry and) pricing, with the worst-handicapped sellers just expecting to recoup their participation costs.

2.2.1.1 Model Specifics

The three risk-neutral agents, named A, B, and C, each start out with one unit of wealth. To obtain (surplus from) trade in equilibrium, we assume the simplest possible mechanism.[2] Each period, one agent is randomly chosen with equal probability 1/3 to be the buyer (BY). The other two agents become potential producers and sellers. There is perfect negative correlation between production and need: the single buyer wants to consume the good but cannot produce it. The two sellers can successfully produce the good but cannot consume it. Each agent knows who is who. The buyer chooses a handicapping bias that may favor a specific seller, and this information is known to the sellers.

Our model must be rigged in a very specific way in order to allow for wealth effects and yet not make equilibrium choices a function of the three parties' wealth. To avoid such path-dependent effects and preserve tractability, our model assumes that all quantities are proportional to the buyer's prevailing wealth. This is not just a convenience assumption, but it allows a seller to gain more surplus when the buyer is wealthier. It is how agents can benefit from increased wealth in their neighborhood. This assumption is also economically reasonable if wealthier buyers like to consume more (expensive) goods. To produce and bring the good to market requires any seller to pay the scaled cost $e{\cdot}W_{t,\mathrm{BY}}$. If BY chooses to buy from a seller, he pays the seller's posted price and receives the good. It's value to the buyer is proportional to his wealth and is $r{\cdot}W_{t,\mathrm{BY}}$. The buyer never pays more than this

[2]The model could be reframed: Instead of this perfect negative correlation between production and consumption, we could assume probabilities of agents succeeding in producing, and wanting to consume with given probabilities. Aside from long and tedious enumerations of different cases (most of which are uninteresting autarchic periods or periods in which exactly one available seller pairs with one available buyer), the simpler model considered here focuses only on an interesting case and can convey the same intended insights.

Figure 2.1: Time Line

Time 0: A and B are neighbors. C is an outsider.

Time t_1: The buyer (BY) for period t is publicly drawn from $\{A, B, C\}$, each with equal probability 1/3. The other two agents become potential sellers.

Time t_2: The buyer announces her handicaps for the two sellers. The buyer's objective is to maximize expected terminal wealth, taking the actions of the other two agents as given.

Time t_3: The two sellers make simultaneous production and pricing decisions. (Sellers cannot observe one another's entry or pricing decision.)

- Each seller makes a decision whether or not to produce. If he does, he pays the entry cost $e \cdot W_{t,\text{BY}}$. The seller may sell a claim to the proceeds of this period's auction to raise the necessary funds.
- Each seller who produces the product posts a price.

The sellers' production and pricing decisions maximize their own expected return over each period,† taking the actions of others as given.

Time t_4: The buyer observes the prices. If there is at least one seller who has produced and priced below the buyer's reservation price, *after* subtracting handicaps, she selects the seller with the lower price. The buyer then pays the posted price to the winning seller, and receives the good he values at $r \cdot W_{t,\text{BY}}$. Unsold goods become worthless.

—: Period t ends. Period wealth is calculated ($\vec{W}_{t+1}$ over $\{A, B, C\}$). With probability $1-\lambda$, the game ends. Otherwise, time t is incremented and the next round commences.

†: Although period-by-period optimization is not necessarily equivalent to terminal wealth optimization, they are isomorphic in the non-reciprocal and fully-reciprocal equilibria discussed below.

reservation price.

The model only makes sense if the buyer values the good more than it costs to bring it to market:

Assumption 1 $r > e > 0$.

Each agent maximizes own expected (cumulative) terminal wealth in a repeated game with uncertain end date. The game ends with probability $1 - \lambda$ after each round. We always assume that

Assumption 2 $\lambda < 3/[3 + r - e]$.

This is necessary to ensure that wealth remains finite in our equilibria.

2.2.1.2 The Zero Constraint

We impose that if a seller cannot afford the participation cost, he cannot offer a product for sale. This may occur because the entry cost is proportional to the buyer's wealth, not the seller's. However, to avoid the dependence of seller behavior on his current wealth level, we also assume that sellers have access to competitive outside investors who break even in expectation. If a seller has decided to enter and is unable to afford the participation cost, he is able to sell a claim to the proceeds of this period's auction. In return, he will receive the expected payoff, which will at least cover the entry cost. The seller then posts a price to maximize the return in this period, following the same strategy as if he had paid the entry cost himself. Because agents are risk neutral, sellers are indifferent to using this financing option in cases where they are able to afford the participation cost.

2.2.2 Endogenous Quantities and Equilibrium Definition

At the beginning of each period, the buyer chooses how much of a premium (s)he is willing to pay for one seller's good over that of another seller. We define three endogenous (strictly positive) handicaps:

1. $\Delta_{t,BC}$, the preference an A buyer has for goods from seller B over seller C in period t;
2. $\Delta_{t,AC}$, the preference a B buyer has for goods from seller A over seller C in period t;
3. $\Delta_{t,AB}$, the preference a C buyer has for goods from seller A over seller B in period t.

After subtracting the handicap from the posted price, buyers choose the cheapest seller. The handicaps are publicly observable.

Equilibrium are the entry decisions as sellers, pricing decisions as sellers, and handicapping decisions as buyers. (We are particularly interested in the equilibrium handicapping choices.) The model has been rigged to keep the optimizing decisions stationary period to period. Buyers choose optimal handicapping decisions to maximize total lifetime wealth. Sellers choose optimal participation and pricing decisions to maximize their return in each period.[3]

Definition 1 (Equilibrium) *An equilibrium is a set*

$$\{ \ (\Delta^*_{t,BC}, \Delta^*_{t,AC}, \Delta^*_{t,AB}), (\tilde{e}^*_{t,A}, \tilde{e}^*_{t,B}, \tilde{e}^*_{t,C}), (\tilde{p}^*_{t,A}, \tilde{p}^*_{t,B}, \tilde{p}^*_{t,C}) \ \}$$

such that, taking the strategies of the other agents as given,

- $\Delta^*_{t,BC}$ *maximizes the expected total wealth of an A buyer;*
- $\Delta^*_{t,AC}$ *maximizes the expected total wealth of a B buyer;*
- $\Delta^*_{t,AB}$ *maximizes the expected total wealth of a C buyer.*

- *When A is a seller,* $\tilde{e}^*_{t,A}$ *is a participation and entry decision* $\{Y, N\}$ *that maximizes his expected return in period* t *(net of entry fee* $e{\cdot}W_{t,BY}$*).*
- *When A has entered, he posts a price* $\tilde{p}^*_{t,A}$ *that maximizes his expected return in period* t*.*
- *When B is a seller,* $\tilde{e}^*_{t,B}$ *is a participation and entry decision* $\{Y, N\}$ *that maximizes his expected return in period* t *(net of entry fee* $e{\cdot}W_{t,BY}$*).*
- *When B has entered, he posts a price* $\tilde{p}^*_{t,B}$ *that maximizes his expected return in period* t*.*

[3] There could be other market structures and equilibria, in which sellers maximize total lifetime wealth. The period-by-period optimization assumption serves to pin down one reasonable entry and price setting mechanism, not the only one. However, the two objectives are isomorphic in the non-reciprocal and fully-reciprocal equilibria discussed below.

- *When C is a seller, $\tilde{e}^*_{t,C}$ is a participation and entry decision $\{Y, N\}$ that maximizes his expected return in period t (net of entry fee $e{\cdot}W_{t,BY}$).*
- *When C has entered, he posts a price $\tilde{p}^*_{t,C}$ that maximizes his expected return in period t.*

The seller's entry $\tilde{e}$ and pricing $\tilde{p}$ are also functions of the (exogenous) parameters, including the identity of the buyer and this buyer's wealth.

The tildes indicate that the optimal entry and pricing choices are (usually) mixed strategies.

Wlog, we will allow only A and B to consider one another as neighbors. From the perspective of C, A and B are interchangeable outsiders. Thus, although C is free to choose a premium for A over B or vice-versa, C will choose not to, because C will never get anything in exchange. This fixes the equilibrium $\Delta^*_{t,AB} = 0$, leaving only two endogenous quantities of interest. Our model's solution is the price premium that an A buyer would be willing to pay for B's good ($\Delta_{t,BC}$) and vice-versa ($\Delta_{t,AC}$), above and beyond what they would be willing to pay for the same good when cheaper from the outside seller C. By symmetry, in equilibrium (but not off-equilibrium), it should be the case that $\Delta^*_t \equiv \Delta^*_{t,BC} = \Delta^*_{t,AC}$, with $\Delta^*_t > 0$.

Also, we will only be considering pure strategies where $\Delta^*_{t,BY}$ is constant over time for each buyer. That is not to say that the handicap is fixed permanently throughout the model. They could choose to deviate in any given period if they so desired, but it will not be rational to do so. There are likely to exist mixed strategies and strategies of cyclic share/don't share behavior, but they do not offer any additional insight.

Note that our equilibrium handicapping strategies are independent of the (earlier) actions by agents. We thus do not consider reciprocity that arises because of sanctions or tit-for-tat strategies. The reciprocity is not achieved or enforced through (off-equilibrium) threats to punish non-compliance (based on earlier behavior). Instead, our model focuses on when *voluntary* reciprocity is sufficient, in which buyers see it in their interests not to defect to the cheapest seller even in the absence of off-equilibrium punishing threats by their neighbors. In equilibrium, A can expect B to continue to buy from A in the future again, even if A were to defect and purchase from C this period instead. Of course, A just chooses not to defect in equilibrium. Moreover, this means that our model can be interpreted either as one in which buyers announce their handicaps once at the outset or they announce their handicaps each period. This is because, in equilibrium, A or B optimally and voluntarily choose a positive handicap Δ_t, anticipating that their neighbor will similarly optimally and voluntarily choose the same positive handicap in the future themselves.

However, it is important that C cannot attempt to bribe A or B to peer with C instead.

2.2.3 Illustration

Suppose $r = 50\%$, $e = 20\%$, $\Delta_{BC} = 10\%$, and $W_{t,A} = \$1$. Then the value to A of buying the good in this period is 50% of his current wealth of \$1, i.e., $r \cdot W_{t,\mathrm{BY}} = \$0.50$. The cost that either seller would have to pay to bring their good to market is equal to 20% of A's current wealth, i.e., $e \cdot W_{t,\mathrm{BY}} = \$0.20$. Finally, A is willing to pay a premium up to 10% of his current wealth for the good of seller B over that of seller C, or $\Delta_{BC} W_{t,A} = \$0.10$. If both sellers have brought their goods to market, and C has posted a price of \$0.30, then A prefers to buy from his neighbor B if B's price is no more than \$0.40. Otherwise, A prefers to buy from C.

Suppose $W_{t,C} = \$0.10$, so C is currently unable to afford the production cost of \$0.20. Because C makes zero profit on average, his expected repayment if he could enter must be equal to \$0.20. However, the actual repayment is uncertain, because the price that C posts is drawn from a mixed strategy. Also, there is some probability that he loses the auction and receives \$0. C therefore sells his claim to this risky payoff, and receives the expected value of \$0.20 in return. He uses this to pay the production cost, bids as normal, and passes the payoff on to the investor. At the end of the day, he will remain exactly where he started, with $W_{t+1,C} = \$0.10$.

Alternatively, suppose $W_{t,B} = \$0.10$, so B is currently unable to afford the production cost of \$0.20. In equilibrium, it will be that B always wishes to enter if he is favored, and if he does enter, he expects to make a profit of $\Delta_{BC} W_{t,A} = \$0.10$. Therefore, the payment he expects to receive from the buyer is equal to \$0.20 + \$0.10 = \$0.30. However, there is uncertainty in his actual payoff, in both the price he will post, and whether or not he wins the auction. In this case, B will sell a claim to this risky payoff and receive the expected value, \$0.30, in return. He will pay \$0.20 of the proceeds towards the entry cost, and keep the other \$0.10. He will bid as normal and pass the payoff on to his investor. B will then end the round with $W_{t+1,B} = \$0.20$.

2.3 Benchmark: Non-Reciprocity

There is always a symmetric equilibrium, in which A and B do not consider one another special.

Theorem 1 *There exists an equilibrium in which no buyer favors any seller:*

$$\Delta^*_{t,BC} = \Delta^*_{t,AC} = \Delta^*_{t,AC} = 0 \; .$$

Each seller follows the same participation and pricing strategy: With probability e/r, this seller does not enter the market. With probability $1-e/r$, this seller pays the participation cost $e{\cdot}W_{t,BY}$ (financing if needed) and posts price $p{\cdot}W_{t,BY}$ according to the cumulative distribution

$$F(\,p\,) = \begin{cases} 0 & \textit{if } p < e \; , \\ \frac{r\cdot(p-e)}{p\cdot(r-e)} & \textit{if } e \leq p \leq r \; , \\ 1 & \textit{otherwise} \; . \end{cases}$$

For the proof, see Appendix 2.9.

Recall that our equilibrium concept assumes that sellers compete perfectly. The participation and price setting mechanisms in this equilibrium serve to leave both sellers with an expectation of zero profit in each and every period. Therefore, conditional on entering, each seller expects to earn back her entry fee $e{\cdot}W_{t,\mathrm{BY}}$. And therefore, the unconditional expected earnings to each seller from the buyer must be $(1-e/r)\cdot e{\cdot}W_{t,\mathrm{BY}}$. Any and all surplus accrues to the buyer.

Note that with probability $(e/r)^2$, neither seller offers the good for sale. In such cases, the buyer, and therefore the system, loses the surplus from trade. In two cases, each with equal probability $e/r \cdot (1 - e/r)$, exactly one of the two sellers offers the good and the buyer pays an expected price of

$$\mathbb{E}(p_t \,|\, 1 \text{ seller}) = \left[\left(\frac{r \cdot e}{r - e} \right) \cdot \log \left(\frac{r}{e} \right) \right] \cdot W_{t,A} \ .$$

Finally, with probability $(1 - e/r)^2$, both sellers enter, and the buyer purchases the good for the (lower) expected price of

$$\mathbb{E}(p_t \,|\, 2 \text{ sellers}) = \left\{ \frac{2 \cdot e \cdot r \cdot [r - e - e \cdot \log(r/e)]}{(r - e)^2} \right\} \cdot W_{t,A} \ ,$$

but there is loss of surplus because two agents have incurred entry costs, when only one was socially productive.

The three identical agents expect to earn over the entire (random) number of periods expected surpluses only when they are the buyers, amounting to

$$\mathbb{E}\,[W_A]^* = \mathbb{E}\,[W_B]^* = \mathbb{E}\,[W_C]^* = \frac{(1 - \lambda) \cdot [3r + (r - e)^2]}{3r - \lambda \cdot [3r + (r - e)^2]} \ .$$

There are two kinds of efficiency losses in this equilibrium. First, it may be the case that both sellers randomly choose not to produce. In this case, the buyer loses surplus. Second, it may be the case that both sellers randomly choose to produce. In this case, the sellers have incurred two participation fees.

2.4 Full Voluntary Reciprocity

It is a more interesting equilibrium when A and B endogenously prefer one another to the outsider C. This equilibrium is viable only if they expect the game to continue long enough.

Theorem 2 *When $\lambda \geq 3/(3 + r)$, there exists an equilibrium in which A and B always buy only from one another,*

$$\Delta^* \equiv \Delta^*_{t,BC} = \Delta^*_{t,AC} = r - e \ .$$

When A or B is the buyer, $e^*_A = e^*_B = 1$ and $e^*_C = 0$. An A or B buyer pays the posted $r{\cdot}W_{t,BY}$ for the good to the other (the neighbor). When C is the buyer, her $\Delta^*_{t,AB} = 0$ and all three agents behave as they do in Theorem 1.

For the proof, see Appendix 2.9.

Although A and B cannot contract, they can rely on their mutual future self-interest. When either is the buyer, the other always enters as the seller and C does not. The buyer pays his entire surplus to the seller, $r{\cdot}W_{t,\mathrm{BY}}$, knowing that when it will be the seller's turn to buy, the seller will return to him.

Note that wealth effects enter our model through scaling of transactions by the buyer's wealth. Each agent must have a reasonable expectation of future benefits from selling transactions, in which the now-buyer will be better off if the local now-seller soon-buyer were to have more wealth. While uncooperatively setting their own Δ, each buyer recognizes that the higher price has an immediate wealth cost to herself now (paying more), but a (multi-period probabilistic) benefit in the future (when selling, the future buyers will be paying more). The equilibrium is maintained, because the latter effect is strong enough in the full-reciprocity equilibrium to outweigh the former, despite the latter's probabilistic nature. This intuition also suggests that the model is plausible only for small neighborhoods. With larger neighborhoods, the return probability (before the world ends) declines too much to make paying the higher price upfront worthwhile.

The comparative statics are simple. By definition, a higher reservation price r and longer-lasting world (λ) means more surplus. Participation e is costly and means less surplus.

The expected terminal wealth of the three agents in this equilibrium is

$$\mathbb{E}\left[W_A\right]^* = \mathbb{E}\left[W_B\right]^* = \frac{(1-\lambda)\cdot(r-e+3)}{3-\lambda\cdot(r-e+3)}\ ,$$

$$\mathbb{E}\left[W_C\right]^* = \frac{(1-\lambda)\cdot[3r+(r-e)^2]}{3r-\lambda\cdot[3r+(r-e)^2]}\ .$$

2.4.1 Comparing Equilibria

Inspection shows that

Theorem 3 *The surplus of the reciprocal over the non-reciprocal equilibrium is*

$$2 \cdot (\mathbb{E}[W_A] - \mathbb{E}[W_C]) = \frac{6\,e \cdot (1-\lambda) \cdot (r-e)}{[3 - \lambda \cdot (3 + r - e)] \cdot \{3r - \lambda\,[(r-e)^2 + 3r]\}} \,.$$

In this full-reciprocity equilibrium, monopoly is socially better than the competition in the no-reciprocity equilibrium.[4] The high price produces no social loss and competition from C would produce no social gain (through the resulting lower price or higher quantities). Exactly one seller, A or B, now chooses to be in the market. This removes the cases in which (1) an A or B buyer find themselves without any potential seller (if both potential sellers had randomly chosen not to appear); and (2) a C never chooses to be in the market (which removes the case in which both sellers incur the entry fee).

The fact that there is a social gain to reciprocity is critically important to this equilibrium. If the total surplus were fixed, and the buyer, A, were simply deciding whether to keep it or pass it to B, she would always keep it. Even under the strongest reasonable wealth effects, the best outcome is that B passes the entire surplus she was given back in the subsequent period. However, due to time discounting, this will be strictly inferior to A keeping the surplus to begin with. It is only because A's options are to keep a small surplus, or pass a *larger* surplus to B, that she is willing to share.

This raises questions whether the social benefits of monopoly are likely to generalize. In smaller markets, monopoly is relatively better than competition, because of reduced duplication and fewer temporary shortages. In larger markets, monopoly is relatively worse than competition, because the law-of-large-numbers is better at balancing demand and supply.

[4]The social advantage of reciprocity—the reduction of duplicate production and the reduction of production failures, brought about by self-sustaining cooperative reciprocity—was not deliberately engineered into the model by us. On reflection, it is of course quite natural. Nevertheless, although the counterveiling economic forces seem intuitively robust, one specific outcome that the outsider C does not suffer seems model-specific.

In a more general model, this would have to be weighed against any inefficient reduction in demand caused by higher neighborly prices.

This equilibrium is near first-best. It is as good as A and B could have contracted for. That is, there is minimal waste even though neither A nor B had to concern themselves with enforcement. The only waste occurs when C is the buyer and A and B may or may not produce.

Theorem 4 *Although C is worse off than A or B in the full-reciprocity equilibrium, the neighbors A and B are better off and the outsider C is no worse off than they are in the no-reciprocity equilibrium in which A and B do not favor one another.*

In our specific model, C does not suffer from discrimination against her compared to the no-reciprocity equilibrium. This is because the economy gains social surplus from coordination at the same rate as C is closed out. This result is unlikely to exactly hold more generically, although the two forces will remain at play.

Because C is no worse off, the comparative statics for when the full-reciprocal equilibrium is better than the no-reciprocal equilibrium are also the same. The reciprocal equilibrium is better when r and λ are higher, and e is lower.

2.5 Partial Voluntary Reciprocity

There are also equilibria in which A and B favor one another, but not so much as to exclude the outsider C from entering the market as a seller at least occasionally:

Theorem 5 *When $\lambda > 3/(3+r)$, there exists an equilibrium, in which the two neighbors A and B reciprocally favor one another modestly but not to the full exclusion of C as a seller,*

$$\Delta^* \equiv \Delta^*_{t,BC} = \Delta^*_{t,AC} = \frac{(r-e)\cdot[\lambda\cdot(r-e)^2 + 3r\cdot(\lambda-1)]}{e\cdot\lambda\cdot(e-2r)} \; .$$

When A or B are buying, their neighbor always produces and randomizes posted prices

$p{\cdot}W_{t,BY}$ according to the CDF

$$F_A(\,p\,) = F_B(\,p\,) = \begin{cases} 0 & \textit{if}\; p \le e + \Delta^* \;, \\ \frac{p-e-\Delta^*}{p-\Delta^*} & \textit{if}\; e + \Delta^* < p < r \;, \\ 1 & \textit{otherwise}\;. \end{cases}$$

(Note that there is a $e/(r-\Delta^*)$ probability mass point at $f(\,r\,)$.)

C produces with probability $1-(e+\Delta)/r$; and, if producing, posts a randomized price $p{\cdot}W_{t,BY}$ according to the CDF

$$F_C(\,p\,) = \begin{cases} 0 & \textit{if}\; p \le e \;, \\ \frac{(p-e)\cdot r}{(p-e-\Delta^*)\cdot(p+\Delta^*)} & \textit{if}\; e < p \le r - \Delta^* \;, \\ 1 & \textit{otherwise}\;. \end{cases}$$

When C is the buyer, he pays no premium $\Delta^*_{t,AB} = 0$. Thus, A and B make entry and offer price decisions the same as they would in Theorem 1.

For the proof, see Appendix 2.9.

The expected terminal wealth of the three agents is

$$\mathbb{E}\,[W_A]^* = \mathbb{E}\,[W_B]^* = \frac{(1-\lambda)\cdot\{3\cdot(r-e)+\lambda\cdot[(r-e)^2+3r]\}}{\lambda\cdot\{3r-\lambda\,[(r-e)^2+3r]\}}\;,$$

$$\mathbb{E}\,[W_C]^* = \frac{(1-\lambda)\cdot[3r+(r-e)^2]}{3r-\lambda\cdot[3r+(r-e)^2]}\;.$$

In this equilibrium, it is never the case that an A or B buyer does not find a seller (as in the non-reciprocal equilibrium), but there is possible duplication: C occasionally incurs an entry cost—even though either A or B have also incurred one. Thus, A and B's surplus and thus the system surplus in this partial-reciprocal equilibrium is always less than the system surplus in the fully-reciprocal equilibrium. Simply put, A and B display the best possible coordination possible in Theorem 2, and partial reciprocity can only do worse.

2.6 Equilibrium Comparisons

Suppose $r = 50\%$, $e = 20\%$, and $\lambda = 0.87$.

In the non-reciprocal equilibrium of Theorem 1, sellers enter with probability $1 - e/r = 60\%$, and randomize over posting prices in the range $[e, r]W_{t,\mathrm{BY}} = [0.2, 0.5]W_{t,\mathrm{BY}}$. With probability $(e/r)^2 = 16\%$ no seller enters, and the buyer enjoys no surplus. With probability $(1 - e/r)^2 = 36\%$, both sellers participate and thus waste one entry cost $e \cdot W_{t,\mathrm{BY}} = 0.2W_{t,\mathrm{BY}}$. Having started with wealth 1, the expected terminal wealth of all three agents is 1.8.

In the fully reciprocal equilibrium of Theorem 2, when A or B are buying, only their neighbor always enters and sells the good for $r \cdot W_{t,\mathrm{BY}} = 0.5W_{t,\mathrm{BY}}$. When the outsider C is buying, A and B behave as they did above. The expected terminal wealth of A and B is 3.3. The expected terminal wealth of C remains at 1.8.

In the partially reciprocal equilibrium from Theorem 5, A or B are willing to pay a premium of $\Delta^* W_{t,\mathrm{BY}} = 0.25W_{t,\mathrm{BY}}$ to one another. Their neighbor will still always enter, but the outsider C now may also enter with probability $1 - (e + \Delta^*)/r = 10\%$. The neighbor randomizes over posted prices in the range $[e + \Delta^*, r]W_{t,\mathrm{BY}} = [0.45, 0.5]W_{t,\mathrm{BY}}$ with an $e/(r - \Delta^*) = 80\%$ probability mass point at the maximum posted price of $0.5W_{t,\mathrm{BY}}$. When participating, the outsider C randomizes over post prices in the range $[e, r - \Delta^*]W_{t,\mathrm{BY}} = [0.2, 0.25]W_{t,\mathrm{BY}}$. The probability that B wins is 91%, and the average price when he does is $0.49W_{t,\mathrm{BY}}$. Conversely, C wins with probability 9% at an average price of $0.22W_{t,\mathrm{BY}}$. The expected terminal wealth for A and B is now 2.9. C again remains at 1.8.

The figures illustrate the tradeoffs and equilbria.

Figure 2.2 shows the best response functions of A and B. When $\lambda > 3/(3 + r)$, there is always one equilibrium in which A and B ignore one another, one equilibrium in which they fully discriminate in favor of one another (to the exclusion of the outsider C), and one equilibrium in which they favor one another, but not extremely so.

Figure 2.3 shows the expected wealth of A, given various choices of A's handicaps for three interesting B handicaps. The plots are of the compounded effect of playing the given strategy forever. The plot of the resulting wealth from the choice of a single periods Δ_t given Δ^* in the future would look similar, but linear and flatter.

Figure 2.4 shows that if A increases his favoritism towards B, and all agents follow optimal participation and pricing strategies, then C is less likely to win. When C wins, he only does so at a lower and lower expected prices.

Figure 2.5 shows that A loses surplus in any given one period when A handicaps in favor of B. However, at the same time, B gains more than A loses (due to elimination of wasteful duplication). Thus, the total system surplus increases when agents rely on reciprocity. It is this effect that ultimately makes it optimal for A and B to reciprocate voluntarily in equilibrium.

The remaining four figures show the equilibrium comparative statics, given our three parameters e, p, and λ.

2.7 Discussion

2.7.1 Period-by-Period Seller Optimization

In our model, the buyer maximizes expected terminal wealth, while the sellers maximize current period wealth. Ideally, sellers would also maximize expected terminal wealth. A seller who is interested in maximizing his expected terminal wealth must also take into consideration how his actions affect the wealth of the buyer and the other seller in subsequent periods. Unfortunately, this problem is generally intractable. Instead, we assume sellers have the simpler objective of myopically maximizing the gain over the given period. Interestingly, the problem becomes time-separable/additive in the no-reciprocity and full-reciprocity cases, i.e., our short-term period-by-period objective for the seller also satisfies the long-term objective.

While we have not solved such a case, we can still theorize how it might look. Suppose the agents are playing our current partial reciprocity game, which maximized per-period gains. A is buying and willing to pay a premium to B. In the short-term sense, B is indifferent to bidding any price, as the expected profit is equal to $\Delta_{t,BC} \cdot W_{t,A}$ regardless. If B bids higher, B wins less often, but gets more if she wins and vice versa. However, if B were to always bid the minimum, $(e + \Delta_{t,BC}) \cdot W_{t,A}$, B would always win. B's profit is still $\Delta_{t,BC} \cdot W_{t,A}$, but now A's profit is $(r - e - \Delta_{t,BC}) \cdot W_{t,A} > (r - e)(r - e - \Delta_{t,BC}) \cdot W_{t,A}/r$, and C loses money. B doesn't care about C's wealth since she doesn't share with B, but B does benefit from future increases to A's wealth since she does share with B. As B increases her bid, more money goes to C, and less goes to A. This isn't a problem with full-reciprocity, as B only ever posts one price. It also isn't a problem with non-reciprocity, since B doesn't care about the others' wealths.

If our goal is to find the bidding strategies that maintain indifference with long-term objectives, we must decrease the attractiveness for B's low bid or increase the attractiveness of B's higher bid. We can't penalize B for her minimum bid since there is no room for C to underbid. We instead need to increase the value today of higher bids to offset the declining long-term value. To do this, C would have to bid less aggressively (i.e. higher prices on average or lower entry probability). Solving for this bidding distribution function is hard, as we need to equalize the long-term payoff values, which compound the effect of this distribution function infinitely. However, this is a promising point for future investigation.

2.7.2 Alternative Pricing Mechanisms

Our model had only three agents who posted prices, and only one single buyer in each period. We determined market participation and pricing through a period-by-period zero-profit condition: Sellers would enter only if they did not expect to lose money. In equilibrium, a non-handicapped seller would always earn exactly this minimal zero profit. Ultimately, we view this as a specificity and/or tractability assumption. We could have constructed a model with similar insights using other (optimizing) participation and pricing mechanisms. For example, we viewed one obvious alternative mechanism, a search market, as less desirable, because our model is about repeated interactions. Presumably, agents would learn over time, e.g., know who their peers are and not have to incur search costs repeatedly. Search economies would further strengthen our model's prediction—that voluntary reciprocity can be sustainable in networks.

The chosen mechanism does have an issue if we wish to add more agents to the model. In the posted price with entry cost model, total welfare decreases as the number of sellers increases (Lang and Rosenthal (1991)). Because all the sellers break even in expectation, average prices increase and total welfare declines as more sellers must recoup their wasted production costs. While each seller enters less often, this does not outweigh the increase in duplicated effort. This effect can be reversed with heterogeneous or uncertain entry costs (Thomas (2002) and Kaplan and Sela (2003)). However, these mechanisms would introduce a great deal of complexity to the model, which is unnecessary to illustrate the simple reciprocity mechanism.

2.7.3 Alternative Arrangements and Larger Networks

We assumed that A and B considered each other neighbors and may not have a preference for C over their neighbor. However, we could also construct an equilibrium where the three agents have a circular relationship: A prefers B, B prefers C, and C prefers A. This arrangement would exhibit the same basic forces, where agents are willing to pass the surplus to another, with the understanding that it will make its way back to them in the long term. The set of parameters for which a full-reciprocity equilibrium exists in the circle is necessarily smaller than for the pair. The extra link adds a time delay, and it will take at least one extra round for the wealth to return. Simply due to discounting, greater gains to coordination through reciprocity are necessary for this equilibrium to obtain.

A possible extension would allow for more than three agents in the model. The question of the reasonableness of the price setting mechanism then becomes an issue. In the model as given, where sellers face homogeneous entry costs and the buyer may favor a single seller, additional agents make the gains to reciprocity *stronger*. Favoritism that dissuades sellers from entering prevents greater wasteful production. Therefore, reciprocal groupings become *more* likely in larger networks. While this general result may seem plausible, the result that it comes from preventing welfare destroying competition seems overly stylized. It's possible that other mechanisms, such as payment for increased matching probabilities or decreased search times in a search model may provide similar scaling incentives.

A larger model would also support a wide number of potential reciprocal configurations. If agents are only allowed to favor one other agent, then any configuration of pairs or smaller loops would exhibit the same reciprocity incentives of expecting gifted surplus to return. However, as the size of the group grows, the incentives decline as the expected wait time to reap the benefits increases.

Endogenous reciprocity is thus more plausible among smaller groups, such as families, tribes, villages, cities, or (coordinated) nations, that are embedded and trading in larger networks.

2.7.4 Assumptions—Tractable vs Economic and Extensions

Our model was highly stylized and had to be rigged to prevent the individual wealth holdings to become full state variables. Thus, it is useful to discuss where its assumptions are particularly objectionable and where they are primarily for tractability's sake.

The presence of potential individual surplus to reciprocity is vital. When all agents are indifferent between buying and not buying, and selling and not selling, they have nothing to gain nothing from reciprocity that harms them at least temporarily. The social surplus makes it easier to obtain the individual surplus.

Our model assumes that A and B view each other as neighbors at the outset. This assumption only serves to simplify the possible space of configurations that need to be considered. Agents can naturally and voluntarily (and without enforcement) choose to become closer, even if there are no natural ex-ante advantages to such team-ups. Of course, exogenous aspects, such as geographic proximity, would likely also play important and reinforcing roles in network self-organization.

Allowing for agents to reconfigure who they choose to pair with at any time would both make the model intractable and weaken our results. Since agents' wealth depends on the wealth of their partners in reciprocal equilibria, they would wish to abandon their current trade partner for whoever is currently wealthiest. This would introduce extreme non-linearities, but would also likely weaken the benefits to reciprocity. If there is some chance that the favored seller will abandon the buyer, they buyer will be hesitant to entrust her with the surplus. The fact that A and B's wealth grows faster than C's in the reciprocal equilibrium makes this event probabilistically unlikely.

The restriction to three parties is also an assumption for tractability. A larger network could be constructed to illustrate the same forces of the reciprocal benefits of favoritism. However, a larger network would be better suited to use an alternative participation and pricing mechanism. The "posted price" mechanism was attractive both in portraying a common economic way to transact and in allowing for interesting model pricing with its possible gain of surplus.

The model insights would be the same if we assumed random successful production and needs. Our model can be viewed as ignoring periods in which agents already have what they need (which they could just consume); in which only one other agent ever happens to have produced (in which there is no choice); in which no agent has produced (nothing to trade); and in which no agent needs the product (nothing to trade).

Relatedly, our model is also not realistic in the consequent effect that there are no intrinsic costs to market power (the monopoly). In this sense, our model is akin to one of perfectly discriminating monopoly. If buyers and the system could gain surplus (e.g., downstream) when prices are lower, the incentive to reciprocity would decline. Our model does not build on monopolistic competition, in which sellers offer differentiated but overlapping goods (in free-entry markets), price their goods competitively, and buyers choose the product best suited to their needs. These kinds of models (such as Helpman (1981) and Krugman (1979)) are not built to investigate the kind of tribal reciprocity that our model is built to explain.

2.8 Conclusion

Our model has shown that agents in a small network can be better off if they voluntarily decide to be neighborly. This reciprocity need not always be enforced. The expectation that the surplus given to a neighbor is more likely to come back than the surplus of an outsider, and thereby improve one's own surplus later, can be enough to sustain reciprocity. In our model, the gains ultimately were higher than those achieved with pure marginal cost pricing, because of surplus gains attributable to improved coordination.

Future research can consider larger networks, ex-ante heterogeneity in costs, and alternative participation and price setting mechanisms.

2.9 Appendix: Equilibrium Derivation

The equilibria from Theorems 1, 2, and 5 are each proved as subcases of the following theorem:

Theorem 6 *The following are Subgame Perfect Nash Equilibria of the infinitely repeated game.*

The selection of handicapping preferences are naive constant quantities that do not depend on the actual past actions of the other agents.

1. *(non-reciprocity)* $\Delta^*_{t,BC} = \Delta^*_{t,AC} = \Delta^*_{t,AB} = 0$, $\forall t \geq 0$ *when* $\lambda < \frac{3r}{3r+(r-e)^2}$
2. *(full-reciprocity)* $\Delta^*_{t,BC} = \Delta^*_{t,AC} = r - e$ *and* $\Delta^*_{t,AB} = 0$, $\forall t \geq 0$ *when* $\frac{3}{3+r} \leq \lambda < \frac{3}{3+r-e}$
3. *(partial-reciprocity)* $\Delta^*_{t,BC} = \Delta^*_{t,AC} = \frac{(r-e)\left(\lambda(r-e)^2+3r(\lambda-1)\right)}{e\lambda(e-2r)}$ *and* $\Delta^*_{t,AB} = 0$, $\forall t \geq 0$ *when* $\frac{3}{3+r} \leq \lambda < \frac{3r}{3r+(r-e)^2}$

Where the following strategies are played by the sellers:

<u>Without a Premium:</u> If the buyer is not willing to pay a premium for either sellers' good, both sellers follow the same bidding strategy independently. With probability $\frac{e}{r}$ *they do not enter the market. With probability* $1 - \frac{e}{r}$ *they pay the entry cost* $e{\cdot}W_{t,BY}$ *(financing if needed) and post a price as a fraction of the buyers wealth* ($p{\cdot}W_{t,BY}$) *according to the distribution:*

$$F(p) = \begin{cases} 0 & \text{if } p \leq e \\ \frac{(p-e)r}{p(r-e)} & \text{if } e < p \leq r \\ 1 & \text{otherwise} \end{cases}$$

<u>With a Premium:</u> If the buyer, A for example, is willing to pay a premium for the good of agent B over the good of agent C ($\Delta_{BC} > 0$)*, then the sellers bid as follows. Agent B always enters and pays the cost* $e \cdot W^A_t$ *(financing if needed). Agent C enters the market with probability* $1 - \frac{e+\Delta_{BC}}{r}$ *and pays the entry cost* $e \cdot W^A_t$ *(financing if needed). They place bids as a fraction of of A's wealth according to the distributions*

$$F_B(p) = \begin{cases} 0 & \textit{if } x \leq e + \Delta_{BC} \\ \frac{p-e-\Delta_{BC}}{p-\Delta_{BC}} & \textit{if } e + \Delta_{BC} < p < r \\ 1 & \textit{otherwise} \end{cases}$$

$$F_C(p) = \begin{cases} 0 & \textit{if } x \leq e \\ \frac{(p-e)r}{(r-e-\Delta_{BC})(p+\Delta_{BC})} & \textit{if } e < p \leq r - \Delta_{BC} \\ 1 & \textit{otherwise} \end{cases}$$

Note that there is a mass point in the distribution of agent B at his maximum posted price $r \cdot W_t^A$*.*

Proof

2.9.1 Sellers' Problem

First, we consider the strategy of the sellers who take the handicap of the chosen buyer as given. We assume that their objective is to maximize their per period expected return, not their expected terminal wealth.

The specific entry and bidding strategies for the sellers are derived in Appendices 2.10 and 2.11. The details that are necessary for determining the handicap strategies are the evolution of expected wealth given handicaps Δ.

In what follows, suppose A is chosen as the buyer in period t. The cases with B and C buying are analogous.

Without a Premium ($\Delta_{t,BC} = 0$): The sellers make 0 profit in expectation, so conditional on entering, which happens with probability $1 - \frac{e}{r}$, they expect to make back the entry cost $e \cdot W_{t,A}$. This revenue comes from the buyer, but the entry cost was paid out of the system. The unconditional expected payment each seller gets from the buyer is then $e(1 - \frac{e}{r})W_{t,A}$.

The buyer receives the good that he values at $r \cdot W_{t,A}$ as long as at least one seller offers it for sale. That is, the expected gain is $(1 - \frac{e^2}{r^2})rW_{t,A}$, and the expected cost is $2e(1 - \frac{e}{r})W_{t,A}$.

We can write the expected profit for the buyer as $\frac{(r-e)^2}{r}W_{t,A}$.

With a Premium ($\Delta_{t,BC} > 0$): Seller B always enters, and his expected profit is $\Delta_{t,BC} \cdot W_{t,A}$. Thus, the expected payment he receives from the buyer is $(e + \Delta_{t,BC}) \cdot W_{t,A}$. Agent C makes zero profit in expectation, so conditional on entering, he expects to make back the entry cost. The expected payment C receives from the buyer is $(1 - \frac{e+\Delta_{t,BC}}{r})eW_{t,A}$.

Because B always enters, A is guaranteed to buy a good that he values at $r \cdot W_{t,A}$. His expected cost is $(e + \Delta_{t,BC} + (1 - \frac{e+\Delta_{t,BC}}{r})e)W_{t,A}$. We can write A's expected profit as $\frac{(r-e)(r-e-\Delta_{t,BC})}{r}W_{t,A}$

2.9.2 Wealth Dynamics

We can then write the evolution of expected wealth, given that A is the buyer as:

$$
\begin{aligned}
E_t[W_{t+1,A}|buyer_t = A] &= W_{t,A} + W_{t,A}(r-e)(r-e-\Delta_{t,BC})/r \\
E_t[W_{t+1,B}|buyer_t = A] &= W_{t,B} + W_{t,A}\Delta_{t,BC} \\
E_t[W_{t+1,C}|buyer_t = A] &= W_{t,C}
\end{aligned}
$$

Letting $\vec{W}_t = [W_{t,A}, W_{t,B}, W_{t,C}]^\top$,

$$
\begin{aligned}
\mathbb{E}_t[\vec{W}_{t+1}|buyer_t = A] &= \left(\mathbb{I}_3 + \begin{bmatrix} \frac{(r-e)(r-e-\Delta_{t,BC})}{r} & 0 & 0 \\ \Delta_{t,BC} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \right) \vec{W}_t \\
&= M_A(\Delta_{t,BC})\vec{W}_t
\end{aligned}
$$

Where $M_A(\Delta_{t,BC})$ is equal to the matrix in parentheses.

We see that all agents' wealth levels are weakly increasing in expectation as long as the buyer's wealth is positive. However, because there is uncertainty in who wins the auction, sellers who pay the entry cost and lose will end up losing wealth. To prevent unlucky agents from ending up with negative wealth, we impose that sellers cannot pay the entry cost if it is greater than their current wealth.

Alone, this would cause bidding strategies to be influenced by the wealth of the sellers, not just the wealth of the buyer. If a seller cannot afford the entry cost, his competitor will post the maximum possible bid. This affects strategies in a non-linear manner that makes expected terminal wealth difficult to compute. To address this, we also assume that the sellers have access to risk-neutral outside investors. If a seller is unable to afford the entry cost, they may sell a claim to the proceeds of this periods auction. Investors are risk neutral, so sellers receive the expected value in return. Sellers then bid according to the same strategy that maximizes this period's proceeds. Now, the seller's wealth will evolve deterministically in same way as it would have in expectation.

Each agent is the seller with equal probability 1/3. If we take the expectation over production/demand outcomes, we can write the evolution of expected wealth as

$$\begin{aligned}\mathbb{E}_t[\vec{W}_{t+1}] &= \left(\mathbb{I}_3 + \frac{1}{3}\begin{bmatrix}\frac{(r-e)(r-e-\Delta_{t,BC})}{r} & \Delta_{t,AC} & \Delta_{t,AB} \\ \Delta_{t,BC} & \frac{(r-e)(r-e-\Delta_{t,AC})}{r} & 0 \\ 0 & 0 & \frac{(r-e)(r-e-\Delta_{t,AB})}{r}\end{bmatrix}\right)\vec{W}_t \qquad (2.1)\\ &= M(\Delta_t)\vec{W}_t\end{aligned}$$

where $\Delta_t = (\Delta_{t,BC}, \Delta_{t,AC}, \Delta_{t,AB})$ and $M(\Delta_t)$ equal the matrix in parentheses.

2.9.3 Buyer's Problem

Suppose at date t, Agent A has been selected as the buyer. He is looking to maximize his expected terminal wealth when the game ends, which occurs at the uncertain date T. Suppose that everyone will be playing their equilibrium strategies $(\Delta^*_{s,BC}, \Delta^*_{s,AC}, \Delta^*_{s,AB})$ in subsequent rounds. If these strategies are constant, and not functions of past actions, we can denote these equilibrium strategies at $\Delta^* = (\Delta^*_{BC}, \Delta^*_{AC}, \Delta^*_{AB})$.

We must consider if, for Agent A, it is profitable to deviate this round and select a $\Delta_{t,BC}$ other than Δ^*_{BC}.

With probability $1-\lambda$, this will be the last round. Since we already know that A is the current buyer, the expected wealth is given by $\mathbb{E}_t[\vec{W}_{t+1}|buyer_t = A] = M_A(\Delta_{t,BC})\vec{W}_t$.

With probability $\lambda(1-\lambda)$, the game will continue one more round past this one. However, in this subsequent round, we do not yet know who the buyer will be. In the following round, agents will resume their equilibrium behavior. Then, conditional on survival, $\mathbb{E}_t[\vec{W}_{t+2}|buyer_t = A] = M(\Delta^*)\,\mathbb{E}_t[\vec{W}_{t+1}|buyer_t = A] = M(\Delta^*)M_A(\Delta_{t,BC})\vec{W}_t$. Continuing on, $\mathbb{E}_t[\vec{W}_{t+s}|buyer_t = A] = M(\Delta^*)^{s-1}M_A(\Delta_{t,BC})\vec{W}_t$.

If we take an expectation over the uncertain end date T, we can write the expected terminal wealth of all agents as

$$\begin{aligned}
\mathbb{E}_t[\vec{W}_T|buyer_t = A] &= \sum_{s=1}^{\infty}(1-\lambda)\lambda^{s-1}\,\mathbb{E}[\vec{W}_{t+s}|buyer_t = A] \\
&= \sum_{s=1}^{\infty}(1-\lambda)\lambda^{s-1}M(\Delta^*)^{s-1}M_A(\Delta_{t,BC})\vec{W}_t \\
&= (1-\lambda)\left(\sum_{r=0}^{\infty}(\lambda M(\Delta^*))^r\right)M_A(\Delta_{t,BC})\vec{W}_t \\
&= (1-\lambda)\left(\mathbb{I}_3 - \lambda M(\Delta^*)\right)^{-1}M_A(\Delta_{t,BC})\vec{W}_t
\end{aligned}$$

The summation term above is a Neumann series. We know that a sufficient condition for convergence is

$$||\lambda M(\Delta^*)||_\infty < 1 \tag{2.2}$$

That is, $\max_i \sum_j \lambda|m_{ij}| < 1$. This depends on the values of Δ^*, so we must assure convergence in each equilibrium we find. If it converges, $\sum_{r=0}^{\infty}(\lambda M(\Delta^*))^r = (\mathbb{I}_3 - \lambda M(\Delta^*))^{-1}$.

This gives us the expected terminal wealth of all agents, however, A is looking only to maximize his own wealth. Lets define his quantity of interest, along those of the other agents, as

$$u_{t,A}(\Delta_{t,BC}|\Delta^*) = [1,0,0](1-\lambda)\left(\mathbb{I}_3 - \lambda M(\Delta^*)\right)^{-1} M_A(\Delta_{t,BC})\vec{W}_t$$
$$u_{t,B}(\Delta_{t,AC}|\Delta^*) = [0,1,0](1-\lambda)\left(\mathbb{I}_3 - \lambda M(\Delta^*)\right)^{-1} M_B(\Delta_{t,AC})\vec{W}_t$$
$$u_{t,C}(\Delta_{t,AB}|\Delta^*) = [0,0,1](1-\lambda)\left(\mathbb{I}_3 - \lambda M(\Delta^*)\right)^{-1} M_C(\Delta_{t,AB})\vec{W}_t$$

where we similarly define

$$M_B(\Delta_{t,AC}) = \begin{bmatrix} 0 & \Delta_{t,AC} & 0 \\ 0 & \frac{(r-e)(r-e-\Delta_{t,AC})}{r} & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad M_C(\Delta_{t,AB}) = \begin{bmatrix} 0 & 0 & \Delta_{t,AB} \\ 0 & 0 & 0 \\ 0 & 0 & \frac{(r-e)(r-e-\Delta_{t,AB})}{r} \end{bmatrix}$$

So, the agents' objectives when they are chosen as the buyer and asked to select a handicap are:

$$A: \quad \max_{\Delta_{t,BC}\in[0,r-e]} u_{t,A}(\Delta_{t,BC}|\Delta^*)$$
$$B: \quad \max_{\Delta_{t,AC}\in[0,r-e]} u_{t,B}(\Delta_{t,AC}|\Delta^*)$$
$$C: \quad \max_{\Delta_{t,AB}\in[0,r-e]} u_{t,C}(\Delta_{t,AB}|\Delta^*)$$

Agent C: It is straightforward to show that $\Delta_{t,AB} = 0$ is optimal. Looking at Equation 2.1 above, we see that increasing $\Delta_{t,AB}$ will cause C to give some of his gains to A, increasing A's wealth next period at the expense of C's. However, none of this surplus ever makes it back to C. C's wealth does not depend on the wealth of A or B at all, as A and B never share any of the surplus back to C. Algebraically,

$$u_{t,C}(\Delta_{t,AB}|\Delta^*) = \frac{3(1-\lambda)(r-e)(r-e-\Delta_{t,AB})}{3r - \lambda(3r + (r-e)(r-e-\Delta^*_{AB}))} W_{t,C}$$

From our parameter assumptions, we have that $\lambda < 1$ and that $r > e > 0$. We require from our convergence restriction on line 2.2 that $\lambda(1 + (r-e)(r-e-\Delta^*_{AB})/(3r)) < 1$. By inspection, $u'_{t,C}(\Delta_{t,AB}|\Delta^*) < 0$ regardless of the choices of agents A and B. It must be true of any equilibrium that $\Delta^*_{BC} = 0$.

Agents A and B: $u_{t,A}(\Delta_{t,BC}|\Delta^*)$ is linear in $\Delta_{t,BC}$, as M is linear in $\Delta_{t,BC}$ and $u_{t,A}$ is linear in M. The slope depends on the choice of agent B, (Δ^*_{AC}). Depending on whether the slope is positive, negative, or zero, agent A will prefer either the maximum ($\Delta_{t,BC} = r - e$), the minimum ($\Delta_{t,BC} = 0$), or be indifferent to all options.

Because we conjecture that the equilibrium strategy is constant, it must turn out that the optimal $\Delta^*_{t,BC} = \Delta^*_{BC}$. So if we conjecture that $\Delta^* = (\Delta^*_{BC}, \Delta^*_{AC}, 0)$, we must have that

$$\Delta^*_{BC} \in \underset{\Delta_{t,BC}\in[0,r-e]}{\operatorname{argmax}} \; u_{t,A}(\Delta_{t,BC}|\Delta^*)$$

$$\Delta^*_{AC} \in \underset{\Delta_{t,AC}\in[0,r-e]}{\operatorname{argmax}} \; u_{t,B}(\Delta_{t,AC}|\Delta^*)$$

Non-reciprocity: $\Delta^* = (0,0,0)$ If B is not sharing any of his surplus with A in the future, then A gains nothing from sharing his surplus with B today. A is not willing to forgo wealth today if he receives nothing in return.

$$u_{t,A}(\Delta_{t,BC}|(0,0,0)) = \frac{3(1-\lambda)(r-e)(r-e-\Delta_{t,BC})}{3r - \lambda(3r + (r-e)^2)} W_{t,A}$$

Here, our convergence restriction from line 2.2 implies that $\lambda < 3r/(3r + (r-e)^2)$. Therefore, the denominator is positive, and $u'_{t,A}(\Delta_{t,BC}|(0,0,0)) < 0$. The optimal $\Delta^*_{t,BC} = 0$, and similarly for agent B.

Full-reciprocity: $\Delta^* = (r-e, r-e, 0)$ Now that B is giving all of his surplus with A, it may be worth it for A to give his surplus to B. The question is whether the increased surplus from B in the future as a perpetuity outweighs the foregone surplus today.

$$u_{t,A}(\Delta_{t,BC}|(r-e,r-e,0)) = \frac{3(1-\lambda)(r-e)\left[\Delta_{t,BC}(\lambda r - 3(1-\lambda)) + 3(1-\lambda)(r-e)\right]}{(3-\lambda(3+r-e))(3(1-\lambda)+\lambda(r-e))r} W_{t,A}$$

Again, our convergence restriction from line 2.2 implies that $\lambda < 3/(3+r-e)$. Therefore, the denominator is positive. We see then that $u'_{t,A}(\Delta_{t,BC}|(r-e,r-e,0)) > 0$ if $\lambda r - 3(1-\lambda) > 0$, or equivalently $\lambda > 3/(r+3)$. If this is the case, then the optimal $\Delta^*_{t,BC} = r-e$, and similarly for agent B.

Partial-reciprocity: $\Delta^* = (\tilde{\Delta}, \tilde{\Delta}, 0)$ where $\tilde{\Delta} = \frac{(r-e)\left(\lambda(r-e)^2+3r(\lambda-1)\right)}{e\lambda(e-2r)}$. Here we have an intermediate case where agents are indifferent to any choice of handicap.

$$u_{t,A}(\Delta_{t,BC}|(\tilde{\Delta}, \tilde{\Delta}, 0)) = \frac{3(1-\lambda)(r-e)^2}{3r(1-\lambda)-\lambda(r-e)^2} W_{t,A}$$

We see that A's choice has no effect on his utility. Therefore, we will let $\Delta^*_{t,BC} = \tilde{\Delta}$ so as to make B indifferent and vice-versa. However, it must be that $0 < \tilde{\Delta}$ and $\tilde{\Delta} < r-e$. Therefore, we require that $\lambda < 3r/(3r+(r-e)^2)$ and $\lambda > 3/(r+3)$ to satisfy the respective requirements. $\lambda < 3r/(3r+(r-e)^2)$ also happens to be our requirement for convergence in this case.

The convergence requirement that $\lambda < \frac{3r}{3r+(r-e)^2}$ is less strict than the requirement that $\lambda < \frac{3}{3+r-e}$. We therefore assume that $\lambda < \frac{3}{3+r-e}$ so any choice of Δ^* will converge.

So far, we considered the expected utility of an agent who already knows he is the buyer at time t. However, we are also interested in his ex-ante expected terminal wealth before a buyer has been chosen. This is given by

$$w_{t,A}(\Delta^*) = [1,0,0](1-\lambda)\left(\mathbb{I}_3 - \lambda M(\Delta^*)\right)^{-1} M(\Delta^*)\vec{W}_t$$
$$w_{t,B}(\Delta^*) = [0,1,0](1-\lambda)\left(\mathbb{I}_3 - \lambda M(\Delta^*)\right)^{-1} M(\Delta^*)\vec{W}_t$$
$$w_{t,C}(\Delta^*) = [0,0,1](1-\lambda)\left(\mathbb{I}_3 - \lambda M(\Delta^*)\right)^{-1} M(\Delta^*)\vec{W}_t$$

2.10 Appendix: Stage Game Without Premiums

A buyer is purchasing one product from one of two sellers with 0 production costs. Sellers submit bids, and the buyer purchases from whoever made the lower bid. The buyer has a reservation price of r. To place a bid, the sellers must pay an entry cost e.

The solution to a more general form of this game with many sellers (but without buyer premiums), was originally derived in Lang and Rosenthal (1991). This is extended by Thomas (2002) to allow for asymmetric entry costs, which reverses price implications of additional sellers, allowing prices to decrease. Kaplan and Sela (2003) achieve similar results with uncertain entry costs.

Symmetric Nash Equilibrium:

- The Entry Game: Each seller randomly decides to enter the market with probability $1 - \frac{e}{r}$, and pays entry cost e if they enter.

- The Pricing Game: Conditional on entering, sellers randomly select a bid according to the distribution below:

$$F(x) = \begin{cases} 0 & \text{if } x \leq e \\ \frac{(x-e)r}{x(r-e)} & \text{if } e < x \leq r \\ 1 & \text{otherwise} \end{cases}$$

  Proceeds conditional on entering are e for both sellers. Thus their expected profit is 0.

Proof:

Consider Seller One:

Seller One's expected revenue from bidding x_1 is given by $\Pi(x_1) = x_1 W_1(x_1) - e$, where $W_1(x_1)$ is Seller One's probability of winning given he bids x_1.

By the FOC, we know that $x_1 W_1'(x_1) + W_1(x_1) = 0$ where:

$$W_1(x_1) = p_2 + (1 - p_2)(1 - F_2(x_1))$$
$$\text{and} \quad W_1'(x_1) = -(1 - p_2) f_2(x_1)$$

Here, p_2 is the probability that Seller Two doesn't enter, and F_2 and f_2 are the CDF and PDF of his bid conditional on entry. Note that the probability of winning is the probability Seller Two doesn't enter, plus the probability that she does enter, but bids higher than x_1 and thus loses. Combining the above formulas:

$$x_1 = -\frac{W_1(x_1)}{w_1(x_1)} = \frac{p_2 + (1 - p_2)(1 - F_2(x_1))}{(1 - p_2) f_2(x_1)}$$
$$\implies f_2(x_1) = \frac{p_2 + (1 - p_2)(1 - F_2(x_1))}{(1 - p_2) x_1}$$
$$\implies f_2(z) = \frac{p_2 + (1 - p_2)(1 - F_2(z))}{(1 - p_2) z} \quad \text{where } z = x_1$$

The solution to this differential equation is given by

$$F_2(z) = \frac{c}{z} + \frac{z}{(1 - p_2) z}$$

If $p_1 > 0$, we know that Seller One must be indifferent between not entering and entering. If he enters, it must be the case that he expects to earn revenue e for any bid, so that his profit after paying the entry cost is 0.

It must be then that the lowest bid is e, as bidding below this would never recoup the entry cost, even if winning were guaranteed. Then $F_2(e) = 0$, so we can solve for the constant c:

$$F_2(z) = \frac{z - e}{(1 - p_2) z}$$

We also know that if Seller One bids the maximum bid, the buyers reserve price r, that he will only win if Seller Two does not enter. Then, the expected profit of bidding r is $\Pi(r) = r p_2 - e = 0$, so $p_2 = \frac{e}{r}$.

$$F_2(z) = \frac{(z - e) r}{(r - e) z}, \quad z \in [e, r]$$

The derivation of Seller One's strategy is identical. He must act so that Seller Two expects to profit zero for any entry decision or bid.

2.11 Appendix: Stage Game With Premiums

A buyer is purchasing one product from one of two sellers with 0 production costs. The buyer is willing to pay a premium Δ for Seller One's product. Sellers submit bids, and if $x_1 < x_2 + \Delta$, the buyer selects Seller One. Conversely, if $x_2 < x_1 - \Delta$, the buyer selects Seller Two. The buyer has a reservation price of r. To place a bid, the sellers must pay a cost e.

Nash Equilibrium: High Type Profits

- The Entry Game:
  - Seller One: Always enters and pays entry cost e.
  - Seller Two: Randomizes to enter the market with probability $1 - \frac{e+\Delta}{r}$, and pays entry cost e.
- The Pricing Game: Conditional on entering, sellers randomly select a bid according to the distributions below:

$$F_1(x_1) = \begin{cases} 0 & \text{if } x \leq e + \Delta \\ \frac{z-e-\Delta}{z-\Delta} & \text{if } e + \Delta < x_1 \leq r \\ 1 & \text{otherwise} \end{cases}$$

$$F_2(x_2) = \begin{cases} 0 & \text{if } x \leq e \\ \frac{(z-e)r}{(r-e-\Delta)(z+\Delta)} & \text{if } e < x_2 \leq r - \Delta \\ 1 & \text{otherwise} \end{cases}$$

Proceeds conditional on entering are $e + \Delta$ for Seller One, and e for Seller Two. Unconditionally, Seller One makes profit Δ while Seller Two makes no profits. Note that $F_1(r) < 1$ so there is a mass point at r.

Proof:

Consider Seller One:

- Seller One's expected revenue from bidding x_1 is given by $\Pi(x_1) = x_1 W_1(x_1) - e$ where $W_1(x_1)$ is Seller One's probability of winning given he bids x_1

- By the FOC, we know that $x_1 w_1(x_1) + W_1(x_1) = 0$ where:

$$W_1(x_1) = p_2 + (1 - p_2)(1 - F_2(x_1 - \Delta))$$
$$W_1'(x_1) = w_1(x_1) = -(1 - p_2) f_2(x_1 - \Delta)$$

  Here, p_2 is the probability that Seller Two doesn't enter, and F_2 and f_2 are the CDF and PDF of his bid conditional on entry. Note that the probability of winning is the probability Seller Two doesn't enter, plus the probability that she does enter, but bids higher than $x_1 - \Delta$ and thus loses.

- Combining the above formulas:

$$x_1 = -\frac{W_1(x_1)}{w_1(x_1)} = \frac{p_2 + (1 - p_2)(1 - F_2(x_1 - \Delta))}{(1 - p_2) f_2(x_1 - \Delta)}$$
$$\implies f_2(x_1 - \Delta) = \frac{p_2 + (1 - p_2)(1 - F_2(x_1 - \Delta))}{(1 - p_2) x_1}$$
$$\implies f_2(z) = \frac{p_2 + (1 - p_2)(1 - F_2(z))}{(1 - p_2)(z + \Delta)} \quad \text{where } z = x_1 - \Delta$$

- The solution to this differential equation is given by

$$F_2(z) = \frac{c}{z + \Delta} + \frac{z}{(1 - p_2)(z + \Delta)}$$

- We impose that $F_2(e) = 0$ as bidding any lower would be below the entry cost and profits would be negative. Therefore, we can solve for the constant c:

$$F_2(z) = \frac{z - e}{(1 - p_2)(z + \Delta)}$$

- We also want the maximum bid to be $r - \Delta$. Bidding any higher than this means that Seller One only wins when Seller Two doesn't enter, so he may as well move any mass in this region up to r. Alternatively, if Seller Two's maximum bid was less than $r - \Delta$, Seller One would similarly be incentivized to move any non-competitive mass to r. It must therefore be that $F_2(r - \Delta) = 1$ and we can solve for $p_2 = \frac{e+\Delta}{r}$.

$$F_2(z) = \frac{(z - e) r}{(r - e - \Delta)(z + \Delta)}, \quad z \in [e, r - \Delta]$$

Consider Seller Two:

- The derivation is similar, but now

$$W_2(x_2) = p_1 + (1 - p_1)(1 - F_1(x_2 + \Delta))$$
$$W_2'(x_2) = w_2(x_2) = -(1 - p_1)f_1(x_2 + \Delta)$$

  as Seller Two only wins if Seller One doesn't show up, or if Seller One bids more than the maximum premium the buyer is willing to pay.

- Now the solution to the ODE is given by:

$$F_1(z) = \frac{c}{z - \Delta} + \frac{z}{(1 - p_1)(z - \Delta)}$$

- Our boundary condition is that $F_1(e + \Delta) = 0$ as bidding below $e + \Delta$ no longer increases Player One's win probability. Therefore:

$$F_1(z) = \frac{z - e - \Delta}{(1 - p_1)(z - \Delta)}$$

- For Seller One, we know that if he bids $e + \Delta$, he will win for sure. The profit of this bid after paying the entry cost e is Δ. Because he must be indifferent between all actions, it cannot be the case that he ever stays out of the market as that would profit 0. It must be the case then that $p_1 = 0$

$$F_1(x_1) = \begin{cases} 0 & \text{if } x \leq e + \Delta \\ \frac{z-e-\Delta}{z-\Delta} & \text{if } e + \Delta < x_1 \leq r \\ 1 & \text{otherwise} \end{cases}$$

If Seller Two bids his maximum bid, it must be the case that $(r - \Delta)W_2(r - \Delta) = e$. That is, his expected profit must be zero at his highest bid. But because Seller One always enters, this means that $(r - \Delta)(1 - F_1(r)) = e$. This means that the probability Seller One bids higher than Seller Two here must be $\frac{e}{r-\Delta}$, but he cannot bid higher than the reserve price r. To fix this, it must be the case that there is a point mass in Seller One's distribution at

r. We impose that Seller Two win all ties here, so that if Seller Two bids his maximum bid $r - \Delta$, he will win with probability $\frac{e}{r-\Delta}$. From Seller One's perspective, the probability that $x_2 = r - \Delta$ and there is a tie is zero. From Seller Two's perspective, if he bids $x_2 = r - \Delta$, it is the same if Seller One bids in the range $[r, \infty)$ and two wins, or Seller Two bids exactly r and two wins the tie.

Figure 2.2: Mutual Best Response

This is a plot of A's best Δ_{BC} response to a given neighbor B's Δ_{AC}, and vice-versa. All agents are assumed to play best participation and pricing strategies, given their Δ's.

The circle is the non-reciprocity equilibrium from Theorem 1. The square is the full-reciprocity equilibrium from Theorem 2. The triangle is the partial-reciprocity equilibrium from Theorem 5.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$.

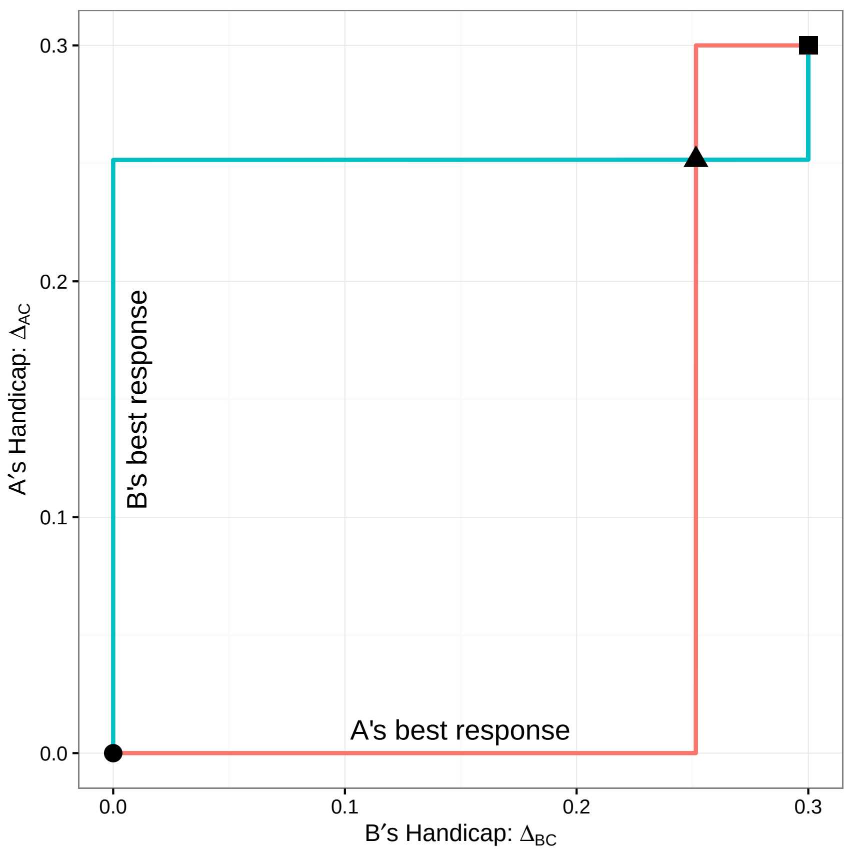

Figure 2.3: A's Payoff as Function of His Imposed Handicap, Given B's Handicap

This figure plots A's expected terminal wealth with respect to his choice of Δ_{BC}, if played forever, for various equilibrium levels of Δ_{AC} (B's chosen handicap). The bottom red and top blue lines show that in the extreme 0 or 0.3 cases of B's handicap, A in turn also prefers an extreme response. However, for a particular Δ_{AC} (middle green line), he is indifferent to all choices of Δ_{BC}. A is not unwilling to share some of his surplus with B if he believes that B is willing to share enough of it (0.25) back in future periods.

All agents are assumed to play best participation and pricing strategies, given their Δ's.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$.

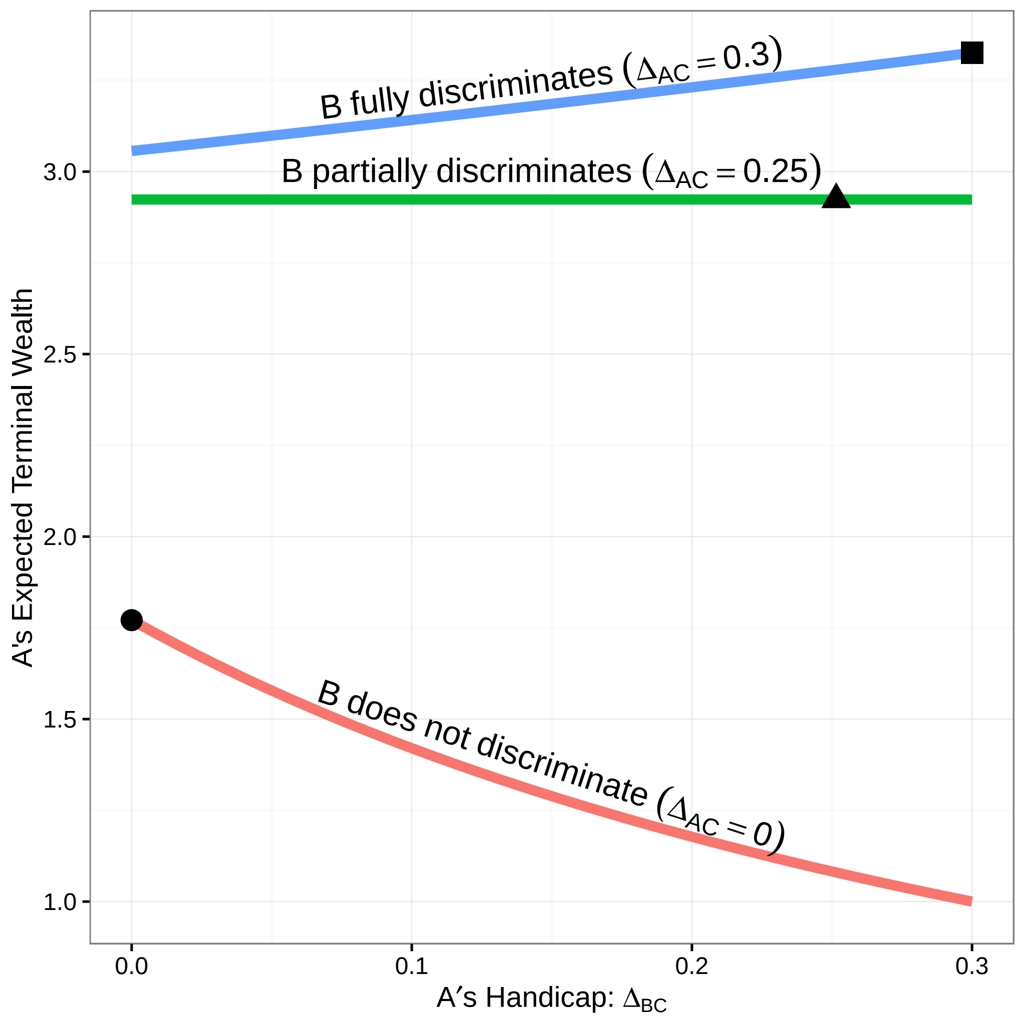

Figure 2.4: Game Dynamics: Winning Probabilities and Pricing as Function of Handicap in Single Period

This figure plots the sellers' per-period win probabilities and winning prices if A were buying, given sellers optimally select their entry and pricing decisions. Here we treat Δ_{BC} as exogenous, and look to illustrate how the sellers respond in their per-period actions.

Left: There is a discrete jump in the probability for $\Delta_{BC} > 0$, because the neighbor B always enters. As Δ_{BC} increases, the outsider C enters less and less often because he is at a disadvantage.

Right: When A is willing to pay a premium, he indeed does so also in equilibrium. For higher handicaps Δ_{BC}, the in-equilibrium winning prices for the outsider C become worse.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$.

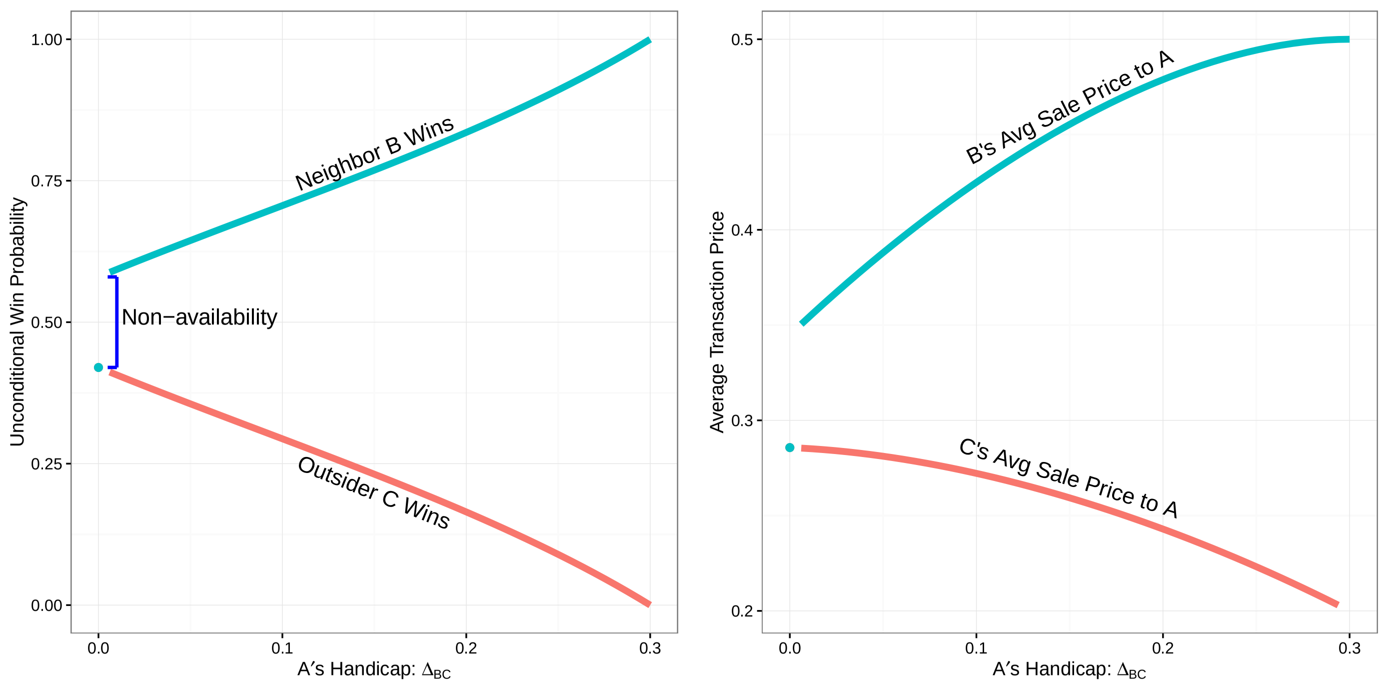

Figure 2.5: Non-Equilibrium Profits in A Single Period

This figure plots the agents' expected profits per-period if A were buying, given sellers B and C optimally select their entry and pricing decisions. Thus, we can look at Δ_{BC} as a choice variable, and plot the returns to both sellers in this one period.

The outsider C always earns zero expected profit. As the handicap that A is willing to pay (to B) increases, A foregoes some of his expected surplus. However, increases in B's surplus are greater than this loss, because the total expected entry costs paid decline.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$.

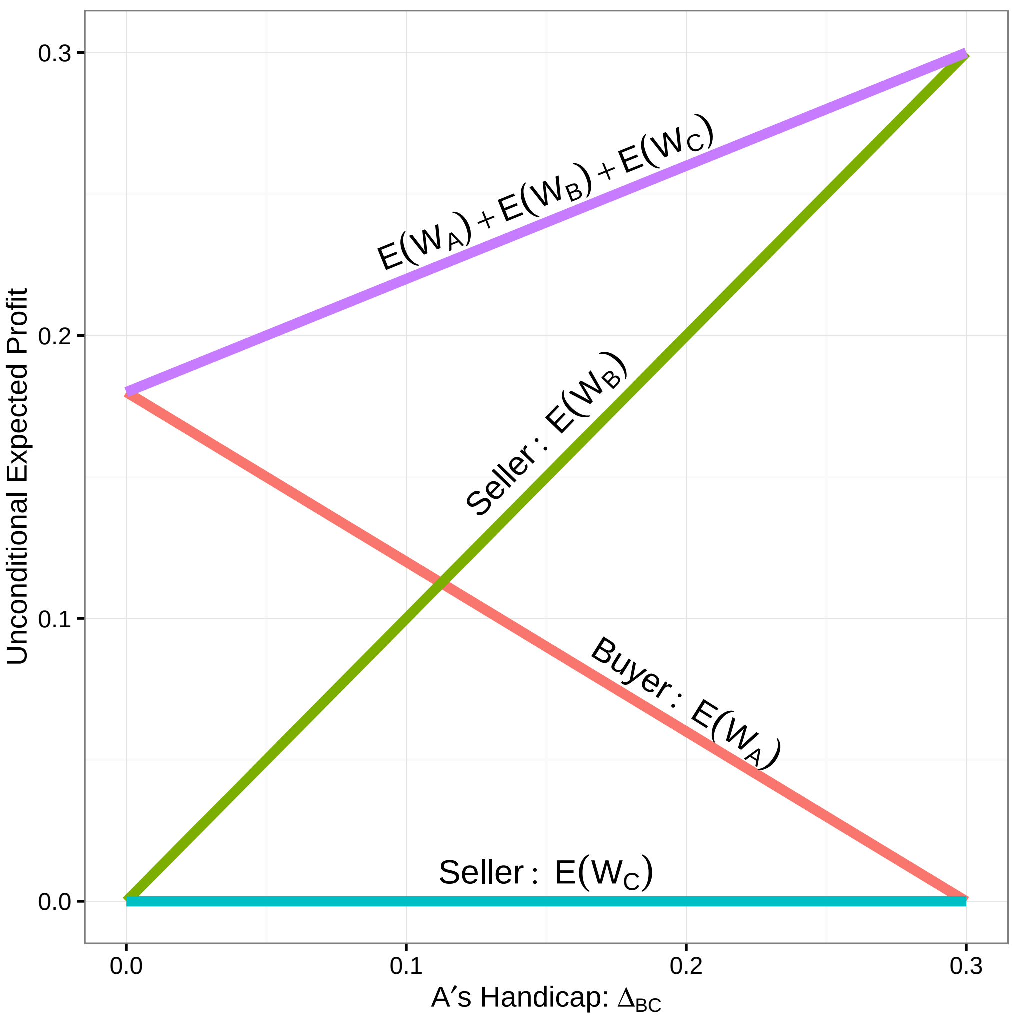

Figure 2.6: Comparative Statics for Equilibrium Handicaps

These plots illustrate changes in the optimal handicap $\Delta^*_{BC} = \Delta^*_{AC}$ for changes in the exogenous parameters. For the non-reciprocal equilibrium, $\Delta^* = 0$, and for the fully-reciprocal equilibrium, $\Delta^* = r - e$. For the partially-reciprocal equilibrium, Δ^* is given in Theorem 5. Note that the latter two equilibria only exist when $\lambda > 3/(r+3)$. It is only when agents value the future sufficiently highly that they are willing to forfeit their surplus as the buyer in exchange for future surplus as the seller. Thus, the reciprocal equilibria begin only at $\lambda = 3/(3+r)$.

Left: Increasing λ decreases the discounting of future wealth. As such, agents are more willing to share their surplus. Optimal Δ in the partially-reciprocal equilibrium must decline to keep the other agent indifferent.

Center: Increasing r increases the total available surplus. When agents share the entire surplus in the fully-reciprocal equilibrium, it is obvious that the surplus shared must increase. However, in the partially-reciprocal equilibrium, the effect may be ambiguous, because agents are balancing the benefit from buyer surplus today with seller surplus tomorrow.

Right: Increasing e decreases the available surplus from trade. When sellers share all the surplus in the reciprocal equilibrium, it is obvious that the surplus shared must decline. The effect on the partially-reciprocal equilibrium is ambiguous.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$, except where a parameter is varying on the X axis.

Length of Game (λ) | Buyer Reservation Price (r) | Seller Participation Cost (e)

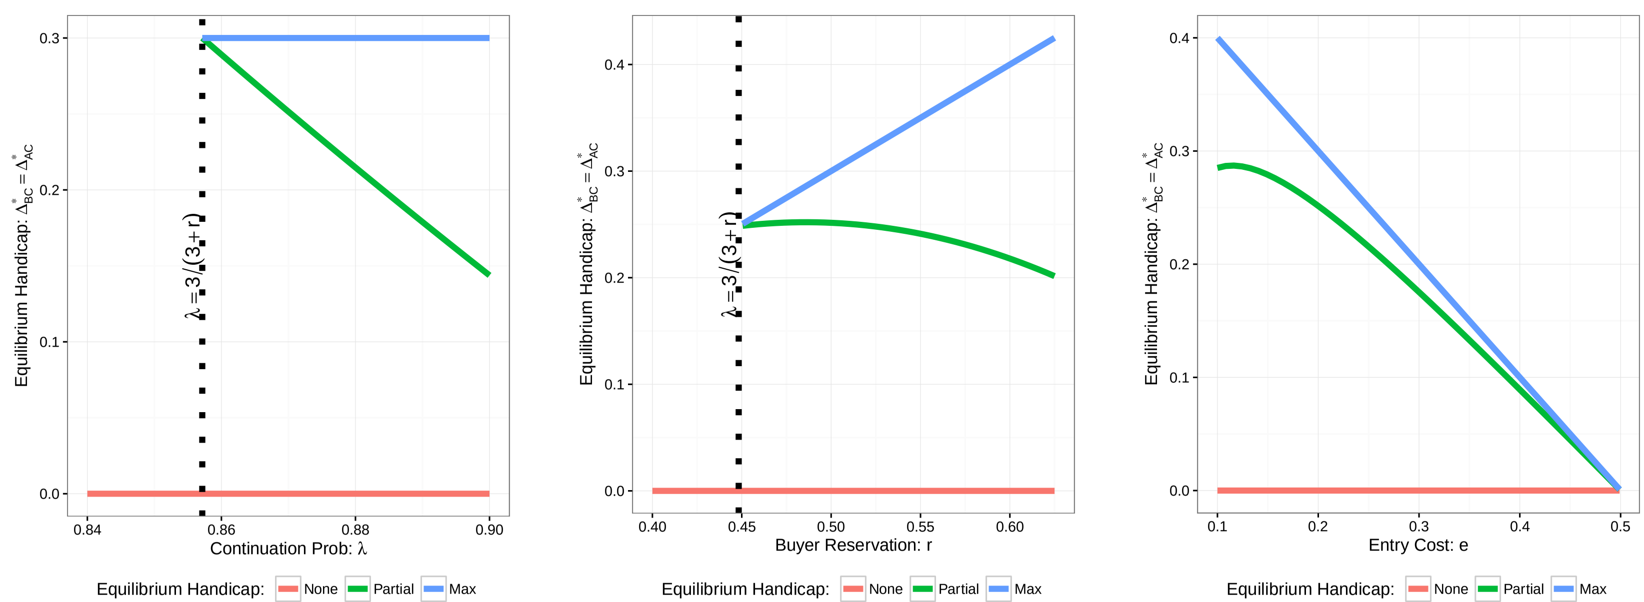

Figure 2.7: Comparative Statics for Equilibrium Wealth

These plots illustrate the changes in the total expected terminal wealth of the economy among the three equilibrium types for changes in the exogenous parameters.

Left: As λ increases, the average game length increases, and the effective discounting of future periods declines. Therefore, the total wealth of the economy grows.

Center: Increasing r increases the value of the asset to the buyer. Total wealth increases with r, but in the reciprocal equilibria that A and B are able to better reap the benefits without duplicated entry costs.

Right: Increasing e increases the costs of bringing goods to market. Total wealth declines as more is paid out of the system in entry costs.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$, except where a parameter is varying on the X axis.

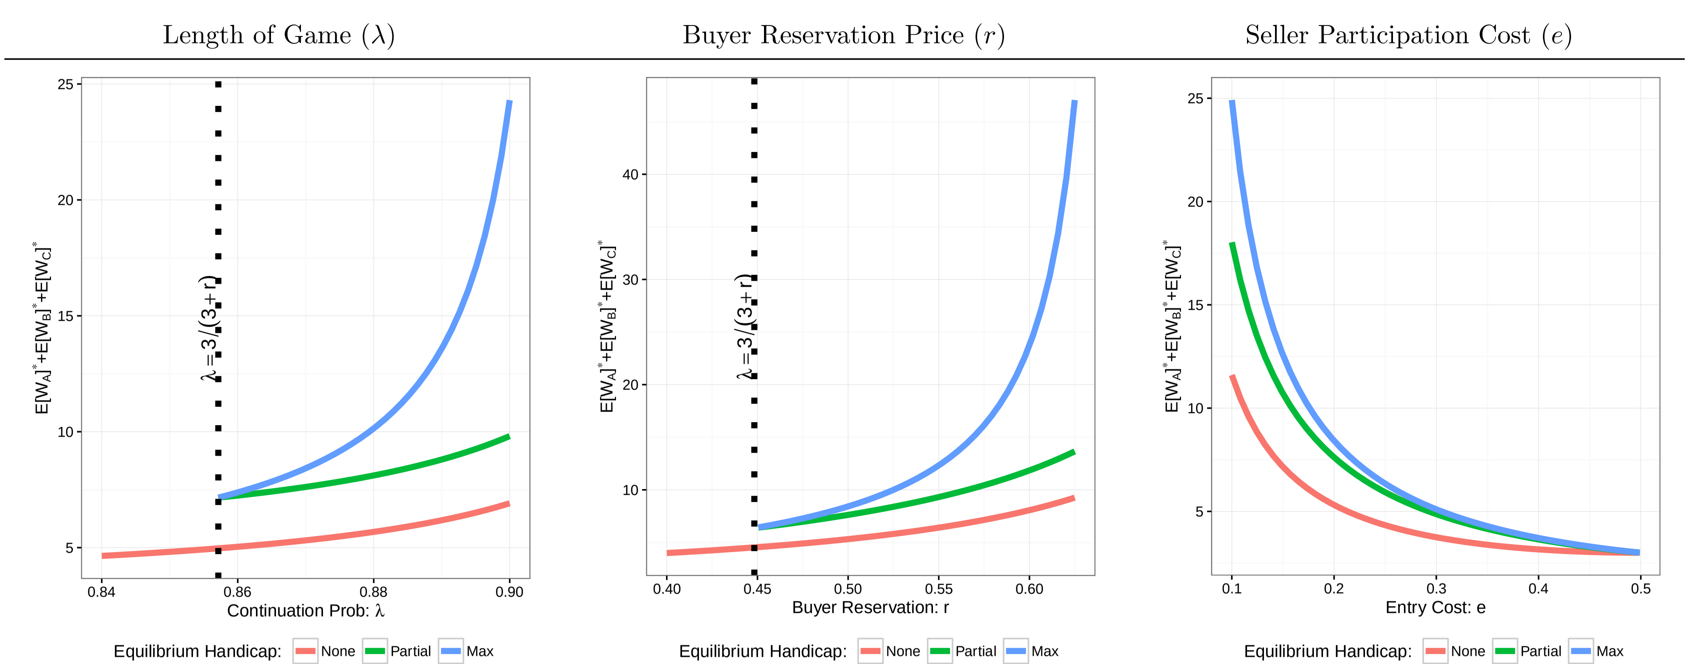

Figure 2.8: Comparative Statics for Neighbor to Neighbor Prices Paid

These plots illustrate the changes in the average price paid by a buyer A or B as a fraction of his wealth to his neighbor among the three equilibrium types for changes in the exogenous parameters. Note that in the full-reciprocity equilibrium, the price is always $r{\cdot}W_{t,\mathrm{BY}}$.

Left: λ does not effect the per period costs and surpluses, so the fully- and non-reciprocal equilibria prices are unchanged. However, the partially-reciprocal equilibrium must balance the trade-off of seller surplus today, compared to buyer surplus tomorrow. As such, increasing λ increases the rewards to sharing, so the price must decline to keep agents indifferent.

Center: Increasing r increases the value of the asset to the buyer. In the non-reciprocal equilibrium, the buyer does not share her surplus so the price is largely unchanged. However, when the seller shares her surplus in the reciprocal equilibria, the price increases.

Right: Increasing e increases the seller's cost to bring the good to market. In the non-reciprocal equilibrium, sellers are competitive so they are only compensated for their cost of entry. Increasing the entry cost must then increase prices. However, in the reciprocal equilibria, the buyer is already paying most of his surplus to the neighbor, so the price is insensitive to the cost.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$, except where a parameter is varying on the X axis.

Length of Game (λ) | Buyer Reservation Price (r) | Seller Participation Cost (e)

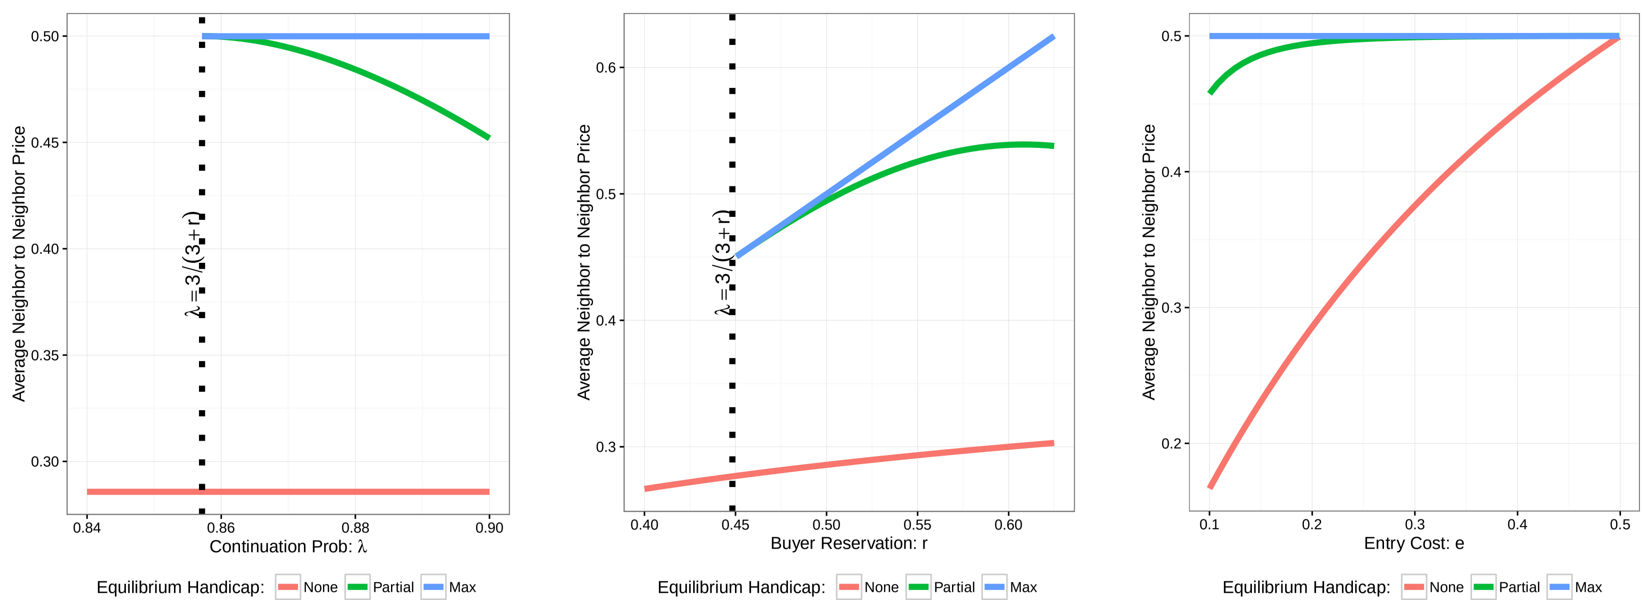

Figure 2.9: Comparative Statics for Probability of Neighbor Purchases

These plots show the probability that a buyer A or B purchases from his neighbor, conditional on him demanding a good. In the non-reciprocal equilibria, the probability that he buys from his neighbor is equal to the probability that he buys from the outsider, but they do not add to one as there is some probability that no good is offered for sale. In the partial and fully-reciprocal equilibria, the buyer always purchases a good, so the probability that he buys from the outsider is complementary.

Left: λ does not effect the per period costs and surpluses, so the full and non-reciprocal equilibria prices are unchanged. However, the partial equilibrium must balance the trade-off of seller surplus today, compared to buyer surplus tomorrow. As such, increasing λ increases the rewards to sharing, so the probability the neighbor wins must decline to keep agents indifferent.

Center: Increasing r increases the probability that the sellers enter the market in the non-reciprocal equilibrium. In the partially-reciprocal equilibrium, it is just the outsider who shows up more often, hence decreasing the probability that the neighbor wins.

Right: Increasing e decreases the probability that sellers enter the market in the non-reciprocal equilibrium. However, in the partially-reciprocal equilibrium, it is just the outsider who shows up less often, hence the neighbor wins more often.

The parameters in this graph are $r = 0.5$, $e = 0.2$, and $\lambda = 0.87$, except where a parameter is varying on the X axis.

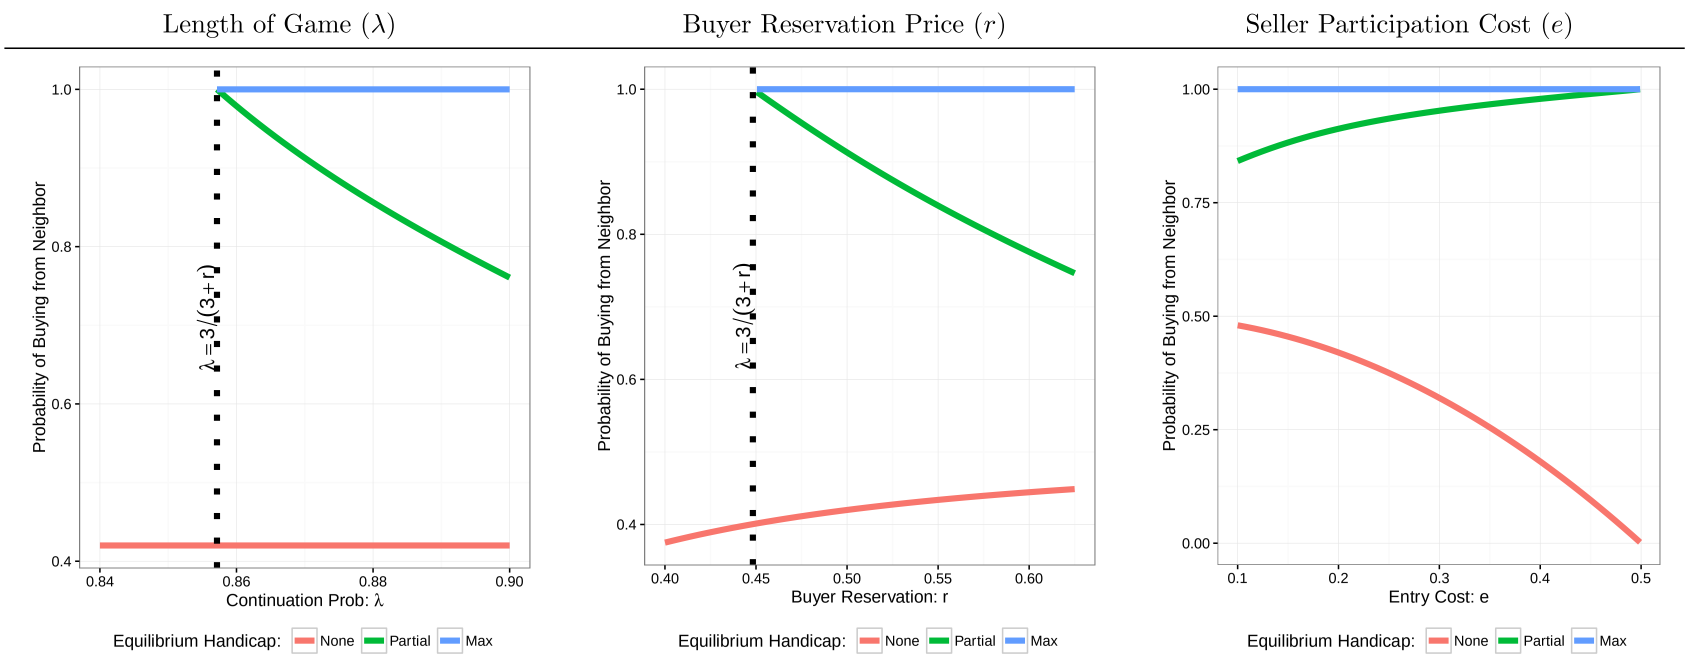

2.12 Bibliography

Axelrod, Robert (1984). *The Evolution of Cooperation*. Basic Books.

Helpman, Elhanan (1981). "International trade in the presence of product differentiation, economies of scale and monopolistic competition: a Chamberlin-Heckscher-Ohlin approach." *Journal of International Economics* 11.3, pp. 305–340.

Kaplan, Todd R and Aner Sela (2003). "Auctions with private entry costs." *SSRN*.

Krugman, Paul R. (1979). "Increasing returns, monopolistic competition, and international trade." *Journal of international Economics* 9 (4), pp. 469–479.

Lang, Kevin and Robert W. Rosenthal (1991). "The Contractors' Game." *The Rand Journal of Economics* 22.3, pp. 329–338.

Thomas, Charles J. (2002). "The effect of asymmetric entry costs on Bertrand competition." *International Journal of Industrial Organization* 20.5, pp. 589–609.

Wilson, David Sloan and E. Sober (1994). "Reintroducing Group Selection to the Human Behavioral Sciences." *Behavioral and Brain Sciences* 17.4, pp. 585–6548.

CHAPTER 3

Asset Redeployability, Liquidation Value, and Endogenous Capital Structure Heterogeneity

with Antonio Bernardo and Ivo Welch

3.1 Introduction

Firms with more leverage are more likely to experience future financial distress. Importantly, their expected costs of bankruptcy are likely to be higher not only when they themselves, but also when their industry peers have taken on more debt. More firms will then want to sell the same types of assets at the same time, and their peer firms—who would otherwise have been the natural asset buyers—become themselves more limited in their capacity to absorb these assets (Shleifer and Vishny (1992)).[1] As a result, the fire-sale discounts relative to fundamental asset values will become steeper. And, therefore, the debt choices of individual firms today, aggregated into industry debt, can themselves influence the asset liquidation values and have anticipative feedback into firms' debt choices in the first place.

Like most earlier literature, in our model, firms choose their capital structures before they learn their profitabilities. Leverage confers direct value benefits, such as signaling benefits, incentive enhancements, or tax shields. However, leverage can lead to distress costs for firms that later experience negative shocks. In the event of default, the creditors must decide whether to liquidate on the one hand, or to reorganize and continue operations on the

[1]Bolton, Santos, and Scheinkman (2011) motivate preferred [industry] purchasers able and willing to pay more than outsiders with adverse selection.

other. If they liquidate, firms receive the prevailing market price for their assets. The assets will then be in the hands of buyers who can presumably put them to better use. If they reorganize, firms keep the assets but may still suffer some impairments, such as direct costs and strained relationships with key stakeholders. A distressed firm is not worth as much as it would have been in the absence of default.

Unlike in most earlier literature, in our model,[2] debt-laden capital-constrained firms are not only more likely to sell but also less likely to buy assets. We assume that all firms are competitive and can anticipate but not internalize the effects of their peers. The mechanism in our model that coordinates their debt choices is the endogenous asset price. For example, suppose that some firms adopt more aggressive debt policies. In the future, this will increase the supply and reduce the demand for liquidated assets, resulting in a lower equilibrium price. In turn, the anticipated lower price creates two motivations for the remaining firms today: (1) they will fear running into financial distress more; and (2), if they reduce their own debt, they will be more likely to enjoy future vulture buying opportunities. Thus, their best response to higher debt by their peers is lower debt for themselves.

The "opportunistic-acquisition" channel can reverse an important implication of models with only the "financial distress" channel. In Williamson (1988) and Harris and Raviv (1990), when assets are more redeployable, firms take on *more* debt because their distress costs will be lower (Benmelech, Garmaise, and Moskowitz (2005)). By contrast, in our model, greater redeployability creates more favorable future buying opportunities and firms may take on *less* debt to take advantage of them. The need of peers to liquidate can create a real growth option in the sense of Myers (1977) or McDonald and Siegel (1986) that can then itself feed back into debt choices and asset prices. No previous model has shown a negative comparative static with respect to redeployability.

Interestingly, when assets are indivisible, a-priori homogeneous firms sometimes split endogenously into two coexisting types who specialize in leverage and role. Some firms

[2]We will discuss the literature in great detail in Section 3.4.3.4.1. Moreover, Table 3.3 shows succinctly how our model's key implications relate to and differ from this earlier literature.

lever up to take advantage of the direct ongoing value benefits of debt—even anticipating distress and having to fire-sell—while other firms maintain conservative capital structures ("dry powder") to take advantage of these anticipated future fire sales (as in the acquisition model of Morellec and Zhdanov (2008)).

Some publicly-traded corporations and industries seem to fit the assumptions of our model. For example, in the shipping industry, where assets are costly and indivisible, Diana Shipping (ticker: DSX) strategically chooses a low-debt conservative capital structure (unlike most other shipping firms) to expand its fleet when ships are liquidated at fire-sale prices.[3] Another natural domain of our model are private companies operating in more local markets. Anecdotal evidence suggests that some local real-estate developers are aggressive, while others wait more patiently for the future fire-sale opportunities in the next downturn. In the context of our model, such heterogeneity can arise more naturally or be amplified for projects such as large local developments (like shopping malls) that are difficult to parcel up.

Our model can also offer further insights. For example, there may be too many or too few asset transfers relative to first-best in our model. And, relevant to the literature on M&A activity, we show that transfer efficiency can be either procyclical or countercyclical, depending on parameters. Thus, for example, any tax policy designed to improve allocational efficiency must be context sensitive. Moreover, our model can also offer predictions on other observables, such as asset transfer quantities and prices, recovery rates and credit spreads, default and liquidation probabilities, and so on.

Our paper also makes a more general point. Most theories of capital structure are about how parameters influence the optimal debt choice. Most empirical tests use normalized leverage, typically dividing it by firm value. When the market value is used, the problem is that not only debt but also firm-value should change with parameters. This matters less

[3]For example, in its 2011 annual report, Diana Shipping stated that its strategy of maintaining a conservative balance sheet allowed it to "seize upon opportunities to deploy our strong cash position to acquire vessels at attractive valuations (p.4)." That year, it used its excess cash to purchase two Panamax dry bulk carriers from distressed sellers at deep fire-sale prices.

when debt and value respond in opposite directions, although what is interpreted as a test on debt could merely be a test on value. It matters more when debt and value respond in the same direction. The empirical metric, debt-to-market-value, then measures merely whether debt or value changes faster. In our specific model, we illustrate this general point by showing how an increase in the direct benefits of debt always increases debt but not always debt-to-value ratios.

Our paper now proceeds as follows: Section 3.2 lays out a basic no-distress model, in which firms with low leverage can later purchase assets from other firms that will turn out to have low productivity. Redeployability favors less leverage, as buyers want the opportunity to purchase poorly performing assets down the line. Debt has an effect only through its influence on this "opportunistic-acquisition" channel. Section 3.3 adds the more recognized "financial-distress" channel. Without the acquisition channel, redeployability always favors more leverage, because sellers can rely on the lesser downside. With both the purchase channel and the distress channel, more asset redeployability at first favors higher leverage (lesser distress costs dominate) and then decreases in leverage (greater acquisition opportunities dominate). Section 3.4 puts the model in perspective and describes its relation to prior research. In particular, it explains why our paper offers the very first model for many of the conjectures in Shleifer and Vishny (1992), and the relation of our model to Gale and Gottardi (2011) and Acharya and Viswanathan (2011). Section 3.5 concludes.

3.2 The Opportunistic-Acquisition Channel

In this section, we introduce a model in which lower leverage allows firms to undertake more acquisitions in the future. In the next section, lower leverage will also reduce expected financial-distress reorganization costs. Table 3.1 summarizes the key variables in our model.

3.2.1 Model Setup and Assumptions

We consider an industry with risk-neutral competitive firms. Each firm has a manager who maximizes the value of the firm. This is not to discount the real-world importance of intra-firm agency conflicts, but to show that our results can obtain even when they are not present.[4] All information is public upon realization, again to show that asymmetric information concerns are not required for our results, not to discount their real-world importance.

Assets and Types: At time 0, each firm owns one indivisible unit of a productive asset.[5] The productivity of this asset is a random variable, denoted $\tilde{v}_i$, whose realization will be publicly observed at time 1. All firms are ex-ante identical and it is common knowledge that their firm type is distributed uniformly on the interval $\tilde{v}_i \in [0, 1]$. After firm productivity is realized at time 1, firms with enough capital (low leverage and high productivity) can acquire assets offered by other firms in the industry at the prevailing endogenous price P. We always assume free disposal, so $P \geq 0$. All assets generate a payoff at time 2, which depends on the holder's realized productivity v_i.

Financing: At time 0, each firm can finance its asset purchase (but not slack excess cash) with equity or debt. The face value of debt is constrained to be $F_i \in [0, 1]$.[6]

All agents are risk-neutral and there is no time discounting, so the expected rate of return on debt is zero. We assume that debt F_i offers immediate net benefits that confer a proportional value $\tau \cdot F_i$. This τ can include the tax benefit of debt (which may or may not be socially valuable), but we have a much broader concept in mind.[7] The parameter τ can

[4]We provide an online appendix in which we model managers that maximize equity and not debt values. Maximizing firm value is the same as maximizing equity value out of default and debt value in default.

[5]At time 0, all firms are identical and we can define their preferred investment amount to be one unit. As we describe below, some firms may wish to purchase liquidated assets at time 1, but these buying firms are then aware of their higher productivity.

[6]An upper limit on F_i ensures that the promised debt payment is never greater than the firm's highest possible cash flow (sans direct benefits). A higher value of F_i would not result in higher proceeds from the debt issuance, because the increment would not be paid. A better assumption would be to impose the upper limit and assess the (tax and other debt) benefits not on the promised but on the expected debt payoff. Unfortunately, this specification forces the model into numerical rather than algebraic solutions.

[7]The model in the text interprets τ broadly as the direct benefits of debt. However, we have solved the model in which τ can represent the tax shield (where taxes also negatively affect firm value), or any combination of tax and non-tax benefits. This requires multiplying our objective functions (except the

reflect the ability of debt to allow financially-constrained firms to take on more productive projects, any positive incentive effects from debt, lower fund-raising costs, and so on; all net of debts' unmodelled costs. This benefit is not dissipated by subsequent events and accrues to the original owners. We show in Appendix 3.8 that all our main results hold when the debt benefit is available to pay creditors and fund acquisitions.

Liquidation: At time 1, after each firm has learned its productivity realization v_i, it can decide whether to sell its asset at the prevailing price P or to continue operations. Because managers' objectives are aligned with their firms', their decision to liquidate or continue is efficient, given their earlier time 0 debt choice. The value from continuing operations is v_i. The liquidation price of the asset is determined by perfectly competitive buyers and sellers. Thus, firm i sells iff $v_i < P$. Although the firm's own debt choice has no influence on the asset's price, each firm knows that the asset price is determined by the collective choices of all firms in the industry.

Acquisition: Although firms can acquire liquidated assets, we assume there is some cost associated with redeployment. This could be because assets need to be customized. Repurposing can require, e.g., moving costs, reprogramming, retraining of workers, and coordination with other complementary assets. In our model, we assume that an asset with productivity v_i to its current owner (firm i) has productivity of $\eta \cdot v_j$ to a potential acquirer (firm j), where $\eta < 1$. Higher values of η imply that assets can be redeployed more easily (at lower cost). In this specification, an asset that transfers from a low-productivity seller i to a high-productivity buyer j enjoys upgraded productivity ($v_j > v_i$), but not to the same extent that it would have had if buyer j had owned it all along. Thus, holding productivity fixed across firms, the asset is also worth more to the current owner than to a potential buyer if both have equal productivity. Taking both firm-specificity and own productivity into account, firms find it worthwhile to buy liquidated assets only if they are sufficiently more productive—acquiring liquidated assets at price P is positive NPV for all firms with

additive $\tau \cdot F_i$ term) by $1 - \tau$. The solutions are appropriately proportional, except that the τ parameter becomes its monotonic transformation $\tau/(1-\tau)$. And with the exception of $\partial V^*/\partial\tau$, which is specifically marked in Table 3.2, all comparative statics remain the same. This is covered in Appendix 3.10.

$v_j > P/\eta$.

As in Shleifer and Vishny (1992), the natural buyers of liquidated assets are other firms in the industry with appropriate expertise. These firms have limited capital, and they are constrained in their ability to acquire the asset at time 1 if they took on too much debt at time 0. Similar limits are also central in other papers, most prominently Duffie (2010). It can be justified, e.g., by *cash-in-the-market* financing (Gale and Gottardi (2015)), where firms are assumed to be unable to raise outside funding on short notice. Our model can go a step further, because it includes one parameter that can help capture at least some cross-sectional or time-series variation in the cash-in-the-market immediacy constraint. Long-term demand curves are more elastic than short-term demand curves. Our parameter η would be lower when assets are shorter-lived, when they require more informed buyers or due diligence (the crisis time relative to the life time of the cash flows), and when they are more difficult to put to use by outsiders. Of course, η also has to encapsulate further real-life aspects, such as how quickly transfer activity would have to take place when aggregate economic and financial conditions are worse. In the extreme allowed in our model, η can approach 1 if industry firms can simply wait out any crisis and search until they can find the nearly perfect buyer.

In our specific model, the only financing available to a firm at time 1 is its internal equity, which is the maximum of zero and $(v_i - F)$. We assume that each firm can only acquire one unit of the liquidated asset at time 1.[8] This reflects limited organizational capacity to take on too many new assets at one time.

Timing: Figure 3.1 illustrates the timing of decisions more precisely.

Objective: At time 0, each firm chooses a debt level F_i to maximize its ex-ante value,

$$V(P, F_i) = \int_0^P P\, dv + \int_P^1 v\, dv + \int_{\min\{P+F_i,1\}}^1 \max\{0, \eta \cdot v - P\}\, dv + \tau \cdot F_i\,. \tag{3.1}$$

[8] The indivisibility assumption is important for the existence of a mixed equilibrium. However, our other qualitative results hold if the asset is divisible and firms can acquire as much of the asset as they can afford with their residual equity $(v_i - F_i)$. See Section 3.2.3.2.4 for more detail.

The first term represents the payoff P if the asset is eventually liquidated ($v_i \leq P$). The second term represents the payoff if the firm chooses to continue operations ($P < v_i \leq 1$). The third term represents the expected surplus if the firm chooses to acquire liquidated assets. The limits of integration recognize that the firm only has sufficient capital to acquire the asset if $v_i \geq P + F_i$ and the integrand ($\max\{0, \eta \cdot v_i - P\}$) recognizes that the firm only acquires the asset if it is positive NPV ($v_i > P/\eta$). The final term represents the immediate benefit of debt.

If the firm's financing constraint is binding (i.e., $1 \geq P + F_i \geq P/\eta$), the expected surplus associated with acquiring assets at time 1 is

$$\int_{P+F_i}^{1} (\eta \cdot v - P)\, dv \;=\; \frac{\eta \cdot [1 - (P + F_i)^2]}{2} \;-\; P \cdot (1 - P - F_i)\,,$$

which is decreasing in the own debt choice F_i. Thus, debt is costly because it reduces future profitable buying opportunities. Furthermore, the surplus is (negative) quadratic in F_i. As the debt level increases, the marginal cost of debt also increases as firms are forced to forgo more and more profitable acquisition opportunities. This is in contrast to the marginal benefits of debt which we have assumed to be linear, leading to the possibility of internal optimal debt levels. Moreover, this cost of debt is increasing when future buying opportunities are of higher quality (i.e., assets are more easily redeployed or the price is lower). When the price of the asset is determined endogenously, as in our model, it will depend partly on how easily the asset can be redeployed. Therefore, the net effect of asset redeployability on equilibrium debt choice is not yet clear.

Because there is no aggregate uncertainty in our model, and we have infinitely many industry participants,[9] firms can anticipate the equilibrium price P at time 0. Therefore, each firm can consider its debt choice in one of three regions, outlined by a marginal cost defined by the right-most integral in (3.1):

1. For low debt, $F_i \leq P/\eta - P$, the marginal cost of debt is zero: increasing debt is not costly because the firm's financing constraint is not binding. Thus, because the

[9] Appendix 3.9 considers an extension with aggregate uncertainty.

marginal benefit of debt is positive (τ), it is always optimal for the firm to increase debt beyond this region.

2. For medium debt, $P/\eta - P < F_i < 1 - P$, the marginal cost is $\eta \cdot (P + F_i) - P$: increasing debt is costly because the firm's financing constraint is now binding, i.e., it may have to forego acquiring positive NPV assets that will be liquidated.

3. For high debt, $F_i \geq 1 - P$, the marginal cost is again zero: the debt is so high that the firm would not be able to finance the acquisition of the asset even if it were to turn out to be the highest productivity, v_i=1. The discontinuous drop in the marginal cost is the result of our indivisibility assumption. At this point the firms cannot afford to purchase an entire asset. If they were allowed to purchase fractional assets, then further debt would still lead to foregone purchases. Increasing debt has no additional costs but additional benefits. Therefore, if the optimal debt is at least $1 - P$, given the τ benefit of debt, it is optimal for such a firm to push its debt to the permitted maximum, here F_i=1.[10]

Together, this means that there are two potential optimal debt levels. One is in the interior region where the marginal benefit is equal to the marginal cost, and one is at the upper boundary where the marginal benefit exceeds the marginal cost but firms have hit the debt constraint. In equilibrium, we will find that firms are sometimes indifferent between these two choices. This means that firms may make different debt choices even if they are identical ex-ante. In particular, firms adopting high-debt strategies (to take advantage of the ongoing debt benefits) will be able to coexist with firms adopting low-debt strategies (to take advantage of future asset buying opportunities at fire-sale prices).

Market Clearing: The equilibrium price for liquidated assets is determined by supply and demand. Because firms may choose different debt strategies, the market clearing price has to be a function of the frequency distribution of firm debt choices. Let $\mathcal{F}(F)$ represent the cumulative distribution function of firms over admissible debt choices $F \in [0, 1]$, i.e., the

[10] All our results hold in a modified model in which the benefits of debt also become available for collateral.

proportion of firms choosing $F_i \leq F$ is given by $\mathcal{F}(\,F\,)$.

Supply: All firms with values $v_i \leq P$ will liquidate their asset regardless of their debt choice.

Therefore, the aggregate supply of the liquidated assets is

$$\int_0^1 \int_0^P 1 \;\; dv\, d\mathcal{F}(F) \;\; (= P) \;. \tag{3.2}$$

Demand: Acquiring one unit of the liquidated asset is positive NPV iff $v_i > P/\eta$. Firms will have sufficient funding to do so iff $v_i \geq P + F_i$, and they will have no demand if they have more debt than $1 - P$. Therefore, the aggregate demand for liquidated assets is

$$\int_0^{1-P} \int_{\max\{P+F, P/\eta\}}^1 1 \;\; dv\, d\mathcal{F}(F) \;. \tag{3.3}$$

3.2.2 Equilibrium

Definition 2 *An equilibrium is a distribution $\mathcal{F}(\,F\,)$ over admissible debt choices $F \in [0,1]$ at time 0 and a price $P \in [0,1]$ for the liquidated asset at time 1, such that*

- *firms act optimally at time 1; and*
- *given a market clearing price P (and their optimal decisions at time 1), firms choose debt F_i to maximize firm value at time 0, according to the distribution $\mathcal{F}(\,F\,)$; and*
- *given the distribution of firm debt choices $\mathcal{F}(\,F\,)$, the price P clears the market for liquidated assets at time 1.*

3.2.3 Solution

Firms are competitive so they take the price of liquidated assets, P, as given. Maximizing ex-ante firm value in equation (1) yields the optimal (interior) debt face value, F^*

$$F^*(\,P\,) = \frac{\tau + (1 - \eta) \cdot P}{\eta} \;,$$

and the maximized firm value of

$$V(\,P, F^*(\,P\,)\,) = \frac{1 + P^2}{2} + \frac{\eta^2 + (P + \tau)^2 - 2 \cdot \eta \cdot (1 + \tau) \cdot P}{2 \cdot \eta} \;.$$

The optimal debt choice is higher when the benefits of debt (τ) are greater and when future acquisition opportunities are worse—when assets are more expensive (P) and more difficult to redeploy (η). However, as we explained above, the equilibrium asset price P^* also depends on the exogenous parameters, so the parameter net effects are yet to be determined.

Together, the equilibrium asset price equates supply, as in (3.2), with demand, as in (3.3); and each firm, given the asset price and its optimal decisions at time 1, chooses debt at time 0 to maximize its value, as in (3.1).

We must also consider the debt choice at the upper boundary, F_i=1, and compare the firm values between the two debt choices. For a given price, the high debt strategy may appear more attractive. However, if all firms choose the maximum leverage, there is no one left to purchase the liquidated assets and the price falls to zero. This makes the interior debt choice more attractive. Mixed strategies may be the only way to balance these forces.

Theorem 7 *In the absence of financial-distress reorganization costs, there exists a unique equilibrium for all parameter values:*

- *If $\tau \leq \eta^2/(3 \cdot \eta + 2)$, there is a pure-strategy equilibrium with price $P^* = (\eta - \tau)/(1 + \eta)$, in which all firms choose $F_L^* = (2 \cdot \tau + 1 - \eta)/(1 + \eta)$.*

- *If $\eta^2/(3 \cdot \eta + 2) < \tau \leq \eta$, there is a mixed-strategy equilibrium with price*

$$P^* = \eta - \tau + \eta \cdot \tau - \sqrt{\eta^2 \cdot \tau^2 + 2 \cdot \eta \cdot \tau \cdot (\eta - \tau)}\,,$$

  in which proportion h^* of firms choose F_H^*=1, and proportion $1 - h^*$ of firms choose F_L^*, where

$$\begin{aligned} F_L^* &= \frac{\tau + (1 - \eta) \cdot P^*}{\eta}, \\ h^* &= \frac{1 - 2 \cdot P^* - F_L^*(\,P^*\,)}{1 - P^* - F_L^*(\,P^*\,)}. \end{aligned}$$

- *If $\eta < \tau \leq 1$, there is a pure-strategy equilibrium with price P^*=0, in which all firms choose F_L^*=1.*

(All proofs are in the appendix.)

If the benefits of debt (τ) are low, all firms choose a low-debt strategy, so that they can maintain financial flexibility to acquire the asset at time 1. For intermediate values of τ, some firms choose a high-debt strategy to take advantage of the immediate benefits of debt, while other firms choose a low-debt strategy to take advantage of future investment opportunities.[11] For high values of τ, the immediate debt benefits outweigh any potential benefit from asset acquisitions, so all firms choose the high-debt strategy. In this case, with no buyers, all assets will end up being discarded rather than being transferred from low-productivity to high-productivity firms.

3.2.4 Implications

A visual perspective can help the intuition. Figure 3.2 plots the comparative statics for heterogeneity h^*.

Type Heterogeneity: This plot shows how heterogeneity in *ex-ante* debt strategies (h^*) arises endogenously. For high redeployability (η) and low debt benefits (τ), all firms choose to operate with very little debt (eager for the opportunity to buy assets from lower productivity firms in the future). For low redeployability and high debt benefits, all firms choose to operate with very high debt (in order to obtain the debt benefits). For intermediate redeployability and debt benefits, ex-ante homogeneous firms naturally divide into two kinds of firms—some pursuing the high-debt operating strategy, others pursuing the lower-debt opportunistic waiting strategy.

This heterogeneity is caused by the indivisibility of the asset.[12] Once a firm has taken on so much debt that it will not be able to purchase the asset, it faces no further marginal

[11] The mixing need not be the same for every firm. The same results obtain in a non-symmetric equilibrium in which h^* firms follow the F_H=1 with certainty, and $1 - h^*$ firms never follow it; or, similarly, any combination of probabilities that lead to an aggregate h^* fraction of firms pursuing F_H=1.

[12] Allen and Gale (1994) discuss divisibility, but their heterogeneity arises from heterogeneity in funding needs.

cost to taking on more debt. If assets were divisible, our comparative statics below would continue to hold, but all firms would act alike. We can thus speculate that heterogeneity in ex-ante strategies increases in asset indivisibility—for example, in real-world situations in which purchases require assuming large pieces (like entire divisions or factories), and not just diversifiable and spreadable small bits and pieces (like retail product inventories).

Implication 1 *When assets are indivisible, ex-ante identical firms may specialize: Low-debt firms coexist with high-debt firms. The region with endogenous heterogeneity is characterized by intermediate levels of redeployability and debt benefits.*

Some of this intuition for leverage and role specialization has also appeared in Morellec and Zhdanov (2008). In their model, there are two potential and strategic acquirers and one target. One potential acquirer decides to specialize in obtaining the tax benefits (with high leverage), while the other specializes in becoming the real acquirer (with low leverage). This is because the equity of a low-debt acquirer does not need to share as much surplus with its own creditors (due to the fact that the debt becomes safer after the acquisition, because the firm becomes larger). The target itself is a third firm, whose value is determined by the competition of the two acquirers. In contrast, in our model, all firms can be acquirers and targets. The leverage of the non-acquiring firms becomes a price-setting component. Despite the obvious similarity, the models also have their differences with respect to heterogeneity. For many parameters, no heterogeneity can emerge in our model. And with many atomistic firms rather than just two, and with firms themselves becoming potential targets, our model can analyze the link between indivisibility and heterogeneity: If there are many firm types and distressed assets are divisible, then all firms would act alike and there would never be heterogeneity (see Section 3.2.4). Moreover, our model's main concern is endogenous distress, itself caused by the very same leverage. We are not aware of any models in the financial distress literature (described further in Table 3.3) that have featured similar endogenous heterogeneity.

Figure 3.3 plots leverage-related comparative statics in this acquisition-channel-only model.

Leverage: The left and middle plots shows that the face value of debt and the value of debt today *decrease* in asset redeployability.[13] This is because, in equilibrium, future buying opportunities are more attractive when assets are more easily redeployed. Hence, firms choose lower debt upfront to enable more opportunistic purchasing in the future. It is this opportunistic-acquisition channel that pushes against the more common intuition that firms take on more debt when their assets are more redeployable because distress costs are lower. Naturally, this implication is robust only to the extent that it characterizes an acquisition constraint. If firms in the industry—broadly defined as firms that are suitable buyers—can purchase liquidating assets regardless of their own leverage (e.g., perhaps because they can raise infinite financing instantly), then this implication is unlikely to hold.

Leverage Ratios: Although debt is unambiguously increasing in τ, the right plots in Figure 3.3 show that this is not true for debt-to-value ratios. The implication of this simple point—that value is also endogenous—is more wide-reaching than just our model. Almost every capital-structure theory has been formulated in terms of debt, while almost every reduced-form empirical capital-structure test has been operationalized in terms of debt-to-value ratios. But with endogenous values, debt-to-value ratios measure primarily the relative speed of the change of debt vis-a-vis the speed of change of value. Thus, empirical test coefficients in naive leverage-ratio regressions may not be translatable into support or rejection of underlying theories.

Implication 2 *Because firm value is also endogenous, comparative statics on debt levels need not be the same as comparative statics on debt-value ratios.*

Peer Effects on Debt Choice: An important aspect of our model is that each firm's debt choice is influenced by its peers via the endogenously determined price of liquidated assets.

[13] With one tiny region exception, which can be seen at the bottom left figure, the discussion applies to both the debt of the low firm (F_L^*) and the debt of the industry ($h^* \cdot 1 + (1 - h^*) \cdot F_L^*$). Our focus is on industry debt, so the discussion omits some trivial tiny-region caveats.

Recall that the optimal (interior) debt choice is

$$F_L^*(\,P\,) = \frac{\tau + (1-\eta)\cdot\eta\cdot P}{\eta}\ ,$$

which is increasing in the price P. The intuition is that future vulture buying opportunities are more attractive when the anticipated asset price is low, so firms have more incentives to reduce debt in order to be more likely to have the financing available to make asset acquisitions. When they conjecture that their peer firms take on more debt, the aggregate demand for the asset declines. The resulting lower equilibrium asset price gives other firms the incentives to reduce debt. This is illustrated in Figure 3.4, which plots the equilibrium price and debt choices as a function of the benefits of debt, τ. (In this example, η=1/2.) For high values of τ, a fraction of firms choose a high-debt strategy (F_H=1), resulting in higher industry debt and a lower asset price than would have obtained if all firms had chosen the low-debt strategy (represented by the dashed-lines). Consequently, firms choosing the low-debt strategy—recognizing that more valuable future buying opportunities will become available—shade their leverage below what would have been optimal if industry debt had been lower.

Implication 3 *Holding parameters constant, with endogenous liquidation values, firms' equilibrium debt choices are negatively influenced by those they conjecture for their peers.*

In real life, peers are likely to have similar parameters for ϕ, η, τ, which would lead them to choose similar capital structures. However, conditional on parameters, higher peer debt gives firms a (marginal) incentive to take less debt, because equilibrium liquidation prices turn lower. However, Leary and Roberts (2014) find evidence even of conditional peer effects, suggesting forces beyond those in our model (such as learning of unknown parameters in industries in which correlated distress is of lesser concern).

3.3 The Distress-Reorganization Channel

The main cost of debt in standard trade-off models like Williamson (1988) and Harris and Raviv (1990) is not debt's constraint on future asset purchases, but its financial-distress cost. Firms that have taken on too much debt will suffer not because they can no longer *buy* when there are fire sales, but because they will have to *sell* when they are in trouble. We now extend our model to show how the two channels work in tandem: the opportunistic-acquisition channel means that debt reduces the demand for liquidated assets, while the financial-distress channel means that debt increases the supply of liquidated assets. Each channel plays the dominant role in some parameter region. Moreover, adding the financial-distress channel makes the model more realistic and adds a wealth of implications.

3.3.1 Setup and Assumptions

The model is similar to the one from the previous section with the following changes:

Impairment: We now assume that there is a dissipative cost when reorganizing-and-continuing in the event of default ($v_i < F_i$). Reorganization here is not necessarily Chapter 11, with its large fixed-cost component, but can also be informal. It seems realistic that the reorganization costs are smaller when the firm is closer to being able to meet its debt obligations.

We specify the distressed reorganization cost to be linear in the shortfall, i.e., $\phi \cdot (F_i - v_i)$. The parameter ϕ represents the losses to a firm's value that are due to being unable to meet pre-agreed debt.[14] The costs could be due to, e.g., direct distractions; damaged relationships with key stakeholders (suppliers, employees, and customers) when the firm is reorganized (Titman (1984));[15] or the residual effects of creditor-manager conflicts (after mitigation by negotiations and side-payments). We model this cost as a proportion of the shortfall rather than a proportion of the total firm value because renegotiation costs are likely to be smaller when the firm is closer to solvency; e.g., is cheap to "buy" bridge funds or leniency from

[14] We focus on the natural case in which $\phi \in [0,1]$. If $\phi \to \infty$, then firms never reorganize. In this region, there is no heterogeneity, but changes in D and D/V are still ambiguous in redeployability η and debt benefits τ. It is still the case that the comparative statics for debt levels can differ from those of debt ratios. The only new comparative static is that Q^* may decrease in η.

[15] For example, Opler and Titman (1994) shows that distressed firms lose market share relative to their conservatively financed peers in industry downturns.

creditors when it concerns one dollar rather than one million dollars.

We always assume limited liability, so firm value under continuation is $\max\{\,0\,,\,v_i-\phi\cdot(F_i-v_i)\,\}$.

Liquidation: At time 1, the manager must decide not only whether to purchase liquidated assets at price P, as in the previous section, but also whether to reorganize in the event of financial distress or liquidate. Financial distress arises when the firm value is below the face value of debt. The firm value is P if it liquidates, and $\max\{\,0\,,\,v_i-\phi\cdot(F_i-v_i)\,\}$ if it continues. For lower firm values v_i, liquidation is better; for higher values, impaired operations is better. Since we assume managers maximize firm value it is straightforward to show that the firm optimally liquidates for all values v_i below a critical value Λ,

$$\Lambda(\,F_i\,)\equiv\frac{P+\phi\cdot F_i}{1+\phi}\;. \tag{3.4}$$

A priori, firms expect to liquidate assets more often when the expected liquidation price (P) is higher and when the relative value from reorganization and continuing operations in distress is lower (i.e., when debt, F_i, is higher or when the reorganization impairment, ϕ, is worse). However, ϕ also has an influence on the equilibrium price, so its net effect is yet to be determined.

Acquisition: As in Section 3.2, the decision whether or not to buy the liquidated asset depends on the transferability of the asset η, the price P, and the firm's capital availability. The asset value to firm i is $\eta\cdot v_i$, so it is positive NPV to acquire the asset iff $v_i>P/\eta$. However, firm i only has sufficient capital to acquire the asset iff $v_i-F_i\geq P$.

Timing: Figure 3.5 illustrates the revised model.

Objective: At time 0, the firm chooses its debt, again taking the expected (and fully anticipated) time 1 price P of liquidating assets as given; and anticipating its own optimal time 1 decisions (a) whether to liquidate or continue operating, and (b) whether to purchase or not purchase other firms' liquidating assets. Therefore, the ex-ante (time 0) value of each firm is

$$V(\, P, F_i \,) = \tag{3.5}$$

$$\tau \cdot F_i + \begin{cases} \displaystyle\int_0^P P \, dv \; + \; \int_P^1 v \, dv + \int_{\min\{P+F_i,1\}}^1 \max\{0, \eta \cdot v - P\} \, dv & \text{if } F_i < P \\ \displaystyle\int_0^{\Lambda(\, F_i \,)} P \, dv \; + \; \int_{\Lambda(\, F_i \,)}^{F_i} [v - \phi \cdot (F_i - v)] \, dv \; + \; \int_{F_i}^1 v \, dv \; + \int_{\min\{P+F_i,1\}}^1 \max\{0, \eta \cdot v - P\} \, dv & \text{if } F_i \geq P \end{cases} \tag{3.6}$$

If the firm takes on less debt than what the asset will be worth, the first row applies and we are back to the case of the previous model. Each firm would know it would operate without possible impairment by distress. If the firm takes on more debt, the second row applies and there are now five terms in the (always-continuous) value objective. The first term is the τ benefit of debt, which accrues immediately.[16] The second term reflects the payoff, P, if the firm is eventually liquidated ($v_i \leq \Lambda(\, F_i \,)$), where $\Lambda(\, F_i \,)$ is given by equation (3.4). The third term represents the payoff to the firm if it is distressed but chooses to reorganize and continue ($v_i \in [\Lambda(\, F_i \,), F_i]$), in which case it receives v_i less the dissipative costs of reorganization $\phi \cdot (F_i - v_i)$. The fourth term is the value of the firm if it is not distressed ($v_i \in [F_i, 1]$) and continues unimpaired. The fifth term represents the expected surplus if the firm acquires liquidated assets. The limits of integration recognize that the firm only has sufficient capital to acquire the asset if $v_i \geq P + F_i$, and the integrand ($\max\{0, \eta \cdot v_i - P\}$) recognizes that the firm only acquires the assets if its NPV is positive given its own type ($v_i > P/\eta$).

Market Clearing: The equilibrium price for liquidated assets is determined by supply and demand:

Supply: As explained above, firms choosing $F_i \leq P$ will liquidate when their realized productivity $v_i \leq P$. Firms choosing $F_i > P$ will liquidate when their realized productivity

[16] In this formulation, the debt benefits cannot be used to stave off liquidation or impairment or to finance the purchase of the asset. However, as already noted above, Appendix 3.8 shows that a model in which firms can do so is isomorphic to the current one. All our principal conclusions continue to hold.

$v_i \leq \Lambda(F_i)$, as described in equation (3.4). Therefore, the aggregate supply of liquidated assets is

$$\int_0^P \int_0^P 1 \; dv \, d\mathcal{F}(F) \; + \; \int_P^1 \int_0^{\Lambda(F)} 1 \; dv \, d\mathcal{F}(\,F\,)\,. \tag{3.7}$$

Demand: Acquiring one unit of the liquidated asset is positive NPV iff $v_i > P/\eta$. Moreover, firms will have sufficient funding to do so iff $v_i \geq P + F_i$. Therefore, the aggregate demand is

$$\int_0^{1-P} \int_{\max\{P+F, P/\eta\}}^{1} 1 \; dv \, d\mathcal{F}(\,F\,)\,. \tag{3.8}$$

3.3.2 Access to Infinite Financing / Eliminating The Acquisition Channel

Before solving the model, it is useful to consider a benchmark in which firms in the industry have infinite access to capital. In this case, the acquisition channel is no longer a constraint. Competition among firms results in an equilibrium with $P^*{=}\eta$, in which (only) the highest-productivity firms ($v_i{=}1$) can purchase *all* assets available for sale. At this high a price, purchasing assets is zero NPV even for the highest-productivity firms and negative NPV for all other firms. Therefore, the acquisition profit terms in both rows in (3.5) drop out. In the first row ($F_i < P$), there is also no disadvantage to raising debt, so firms would always be better off increasing debt and leaving this region. This leaves only the second row for consideration.

Substituting $\Lambda(\,F_i\,) = (P + \phi \cdot F_i)/(1 + \phi)$ into the objective and taking the derivative with respect to F_i yields the first-order condition for the (interior) optimal debt choice. The symmetric pure-strategy equilibrium debt choice is

$$F^* = P^* + (1 + 1/\phi) \cdot \tau = \eta + (1 + 1/\phi) \cdot \tau\,.$$

The optimal debt choice is increasing in the benefits of debt (τ) and asset redeployability (η), and decreasing in the costs of reorganization (ϕ). This is the standard result in earlier literature. In particular, debt increases in asset redeployability, because more redeployable assets have higher liquidation values ($P^*{=}\eta$), thereby reducing distress costs. There is no countervailing cost of debt with unlimited capital—increasing debt never precludes firms from acquiring valuable (high η) assets.[17]

3.3.3 Equilibrium With Acquisitions and Financial Distress Channels

The trade-offs associated with the firm's debt choice in our general model with both the opportunistic-acquisition channel and the distress-reorganization channel depend again on the level of debt F_i vis-a-vis the predictable asset price P. For example, consider the case where $\eta \geq 1/2$. (All cases are derived in the Appendix.) Each firm takes P as given and considers its possible debt choice in one of four distinct regions:

1. In the first region, $F_i \leq P/\eta - P$, the marginal cost of debt is zero. Firm value is described by the first row of equation (3.5): increasing debt does not increase reorganization costs (because the firm will always liquidate in distress) and the firm does not forego asset-acquisition opportunities (because the financing constraint is not binding).

2. In the second region, $P/\eta - P < F_i \leq P$, the marginal cost of debt is $\eta \cdot F_i - P \cdot (1 - \eta)$. Firm value is still described by the first row of equation (3.5): increasing debt still does not increase reorganization costs, but the firm now may forego some positive NPV asset-acquisition opportunities.

3. In the third region, $P < F_i < 1 - P$, the marginal cost of debt is $\eta \cdot F_i - P \cdot (1 - \eta) + (F_i - P) \cdot \phi/(1 + \phi)$. Firm value is now described by the second row of equation (3.5): increasing debt raises the expected reorganization costs *and* results in the firm foregoing some positive NPV buying opportunities.

4. In the fourth region, $1 - P \leq F_i \leq 1$, the marginal cost of debt is $(F_i - P) \cdot \phi/(1 + \phi)$. Firm value is again described by the second row of equation (3.5): the firm's debt is

[17] Substituting F^* into Λ yields the equilibrium liquidation threshold $\Lambda(F^*) = P + \tau = \eta + \tau$. The optimized firm value is

$$V^* \equiv V(F^*) = \frac{1 + (\eta + \tau)^2 + \tau^2/\phi}{2}.$$

now so high that it would *never* be able to buy assets even if it turned out to be the highest productivity type, v_i=1. Therefore, the only remaining marginal cost of debt is the increase in expected reorganization costs.

Importantly, the marginal cost of debt is weakly increasing in F_i over the first three regions, but then jumps down at F_i=$1 - P$ (because the firm can now never afford to purchase the asset), after which it increases again. Consequently, as in our model without reorganization costs, there is again a region with a mixed equilibrium, in which some firms choose low debt and others choose high debt.

Equilibrium requires again that firms make optimal decisions at time 1 (both continuation and asset acquisition); their debt choices at time 0 maximize firm value (3.5), given the anticipated asset price and optimal decisions at time 1; and the market for liquidated assets clears, i.e., supply in equation (3.7) is equal to demand in equation (3.8).

The description of the equilibrium solution for all parameters is very detailed and depends on different parameter regions for the reasons just described. Therefore, for the sake of the exposition, in the following theorem we describe equilibria for a particularly relevant parameter region—when ϕ is modest and η is large, for all values of τ—and leave the full description and proof of the theorem for all parameters for Online Appendix A.

Theorem 8 *Assume (i) $\eta \geq 2/3$ and (ii) $\phi < (3\eta - 2)/(6 - 3\eta)$ and let*

$$\tau_1 = \frac{2(\phi+1)\eta^2 - \phi + \eta\cdot(\phi+1) - \sqrt{(\eta+1)^2\cdot(\phi+1)\cdot(\eta^2\phi+\eta^2-2\phi-\eta\phi)}}{3\eta(\phi+1)+3\phi+2},$$

$$\tau_2 = \frac{(2\eta-1)\cdot(2\eta+\phi+2\eta\phi) - \sqrt{(2\eta-1)^2\cdot(1+\phi)\cdot[4\eta^2(1+\phi)+\eta\phi-2(\eta+\phi)]}}{2+3\phi},$$

$$\begin{aligned}\tau_3 = {} & \frac{2\eta^2(1-\phi^2)+\eta(1+9\phi+\phi^2-5\phi^3)+2\phi+12\phi^2+7\phi^3}{2+14\phi+20\phi^2+9\phi^3+3\eta\cdot(1-\phi)\cdot(1+\phi)^2} \\ & + \frac{\sqrt{(1+\phi)\cdot(1+\eta+5\phi-\eta\phi)^2\cdot[\eta^2(1+\phi)(1+3\phi^2)-3\eta\phi^2(2+\phi)-2\phi^3]}}{2+14\phi+20\phi^2+9\phi^3+3\eta\cdot(1-\phi)\cdot(1+\phi)^2},\end{aligned}$$

$$\tau_4 = \frac{\eta+\phi+\eta\phi}{1+2\phi}.$$

For any set of parameter values that satisfy the above restrictions, there exists a unique equilibrium. The following is a complete characterization of the equilibrium:

- *If* $0 \leq \tau \leq \tau_1$, *there is a pure-strategy equilibrium with price*

$$P^* = \frac{\eta - \tau}{1+\eta} \ ,$$

in which all firms choose

$$F_L^* = \frac{1-\eta+2\tau}{1+\eta} \ .$$

- *If* $\tau_1 < \tau \leq \tau_2$, *there is a mixed-strategy equilibrium with price*

$$\begin{aligned} P^* &= \frac{\phi\eta - (1+\phi)\cdot[\tau - \eta(1+\tau)]}{1+\phi(1+\eta)} \\ &\quad - \frac{\sqrt{\eta(\phi+1)\cdot\{2\tau\cdot[\eta\phi+(\eta-\tau)\cdot(1+\phi)]+\eta\tau^2(\phi+1)-\phi\cdot(1+\tau-\eta)^2\}}}{1+\phi(1+\eta)} \ , \end{aligned}$$

in which fraction h^* *of firms choose* F_H^*=1, *and fraction* $1-h^*$ *choose* F_L^*, *where*

$$\begin{aligned} F_L^* &= \frac{\tau}{\eta} + \frac{(1-\eta)}{\eta}\cdot P^* \ , \\ h^* &= \frac{(1+\phi)\cdot[\eta-\tau-(1+\eta)\cdot P^*]}{\eta\phi+(1+\phi)\cdot(\eta-\tau)-[1+\phi(1+\eta)]\cdot P^*} \ . \end{aligned}$$

- *If* $\tau_2 < \tau \leq \tau_3$, *there is a mixed-strategy equilibrium with price*

$$\begin{aligned} P^* = &\frac{\phi\cdot[1+2\phi(1-\tau)-3\tau]+\eta(1+\phi)\cdot[1+\tau+(2+\tau)\phi]-\tau}{1+(6-3\eta)\cdot(1+\phi)\cdot\phi} \\ &- \frac{\sqrt{(1+\phi)\cdot(\eta+\phi+\eta\phi)\cdot\left\{\begin{array}{l} 3\eta^2\phi(1+\phi)-2[\phi(\tau-1)+\tau]^2 \\ \quad +\eta[\phi(\tau-1)+\tau]\cdot[2+(\tau-1)\phi+\tau] \end{array}\right\}}}{1+(6-3\eta)\cdot(1+\phi)\phi} \ , \end{aligned}$$

in which h^* *firms choose* F_H^*=1, *and* $1-h^*$ *choose* F_L^*, *where*

$$\begin{aligned} F_L^* &= \frac{(1+\phi)\cdot\tau}{\eta+\phi+\eta\phi} + \frac{(1-\eta)\cdot(1+\phi)+\phi}{\eta+\phi+\eta\phi}\cdot P^* \\ h^* &= \frac{(1+\phi)\cdot[\eta+\phi+\eta\phi-(1+2\phi)\tau-(1+\eta+5\phi-\eta\phi)\cdot P^*]}{(1+2\phi)\cdot[\eta+\phi+\eta\phi-(1+\phi)\tau]-(1+5\phi-\eta\phi+5\phi^2-\eta\phi^2)\cdot P^*}. \end{aligned}$$

- *If $\tau_3 < \tau \leq \tau_4$, there is a pure-strategy equilibrium with price*
$$P^* = \frac{\eta + \phi + \eta\phi - \tau \cdot (1 + 2\phi)}{1 + \eta + 5\phi - \eta\phi} \ ,$$
in which all firms choose
$$F_L^* = \frac{1 + 2\phi + \tau + (\tau - \eta) \cdot (1 + \phi)}{1 + \eta + 5\phi - \eta\phi} \ .$$
- *If $\tau_4 < \tau \leq 1$, there is a pure-strategy equilibrium with price P^*=0, in which all firms choose*
$$F_L^* = \min \left\{ 1 \ , \ \frac{\tau(1 + \phi)}{\eta(1 + \phi) + \phi} \right\} \ .$$

3.3.4 Implications

Unlike our model from the previous section, debt is now costly for two reasons: first, it reduces future purchasing opportunities; *and* second, it increases the expected costs of financial distress. The model still has only three parameters—the redeployability of assets (η), which is central to our acquisition channel; the reorganization impairment parameter (ϕ), which is central to our financial distress channel; and a compensating direct benefit of debt (τ). Yet, the model can offer many implications. Of course, it remains too stylized to consider its implications to be either quantitative or universal. Instead, our model should be viewed as suggestive of economic forces in contexts in which both the financial-distress and the opportunistic-acquisition channels are important for firms that can become either sellers or buyers of distressed assets in the future.

This subsection discusses the model's comparative statics. They are summarized in Table 3.2 and illustrated in the graphs that follow. The graphical approach is more intuitive, although the model's implications are also algebraically demonstrable using the closed-form solutions in Theorem 8.

Figure 3.6 shows the proportion of firms choosing maximum debt, firm values, and leverage in the case in which ϕ=0.25. This parameter means that reorganization would consume one quarter of each dollar's *shortfall*. This seems high for large firms, although it is not unreasonable for midsize and smaller firms (Bris, Welch, and Zhu (2006)).

3.3.4.1 Heterogeneity (h^*)

As in the model of the previous section, heterogeneity in ex-ante leverage strategies can arise endogenously because our assets are indivisible. Figure 3.6 shows the now parabolic convex region that separate the set of homogeneous (pure) from the set of heterogeneous (mixed) equilibria. These mixed equilibria occur again when otherwise identical firms infer that the corner solution, with maximum permitted debt of F_H=1, is as good for them as the best interior debt choice (F_L^*). Not surprisingly, mixed equilibria can only occur in regions in which firms want to choose fairly high debt to begin with.

Comparing the heterogeneity when ϕ=0 in Figure 3.2 with its equivalent when ϕ=0.25 in Figure 3.6 shows that reorganization costs ϕ shrink the heterogeneous region. For sufficiently low values of either debt benefits τ or redeployability η, there are now only homogeneous equilibria. Nevertheless, the set of mixed equilibria remains non-trivially large. In detail:

- When the debt benefits τ are low, all firms choose low debt because the benefits of a high-debt strategy are too small to compensate for the foregone investment opportunities. Similarly, when the redeployability η is low, liquidation values are low and again all firms choose low debt because a high-debt strategy results in excessive distress costs

- At some point, with high enough redeployability and debt benefits, some firms can begin to specialize in waiting for acquisition opportunities. Heterogeneous equilibria appear only for intermediate values of τ and high values of η. Thus, the heterogeneous region becomes smaller than it was in Figure 3.2.

- Finally, when the debt benefits become overwhelming, all firms end up choosing high debt and no firm finds it worth waiting for opportunities, even though such firms expect large distress costs.

3.3.4.2 Firm Value

Firm value always decreases monotonically in reorganization costs ϕ.

Firm value also usually increases in redeployability η. However, it can occur in a tiny parameter region (with high redeployability, low reorganization costs, and high benefits)—too small to be even visible in this graph—that firm value can decrease.

The effect of direct debt benefits (τ) on firm value is our *only* comparative-static implication that depends on the source of the direct benefits of debt:

- If the debt advantages are not from taxes but from incentive or information causes (and purely additive), as in the model presented here in our main text, then firm value always increases in τ.

- If the debt gains are from taxes, firm value can still increase or decrease in τ. This is somewhat surprising. As expected, for small tax-rates, taxes reduce the firm value directly (through their multiplicative $1-\tau$ factor on the value part of the objective function). The levered firm merely is less negatively effected by the required tax payments. However, for higher tax rates (about halfway up in the feasible region), equilibrium firm value also increase again in the tax rate. This is partly due to the ability of firms with very low expected values to resell the still-valuable tax credits on the market, and partly due to an equilibrium effect that is determined by the interplay of leverage and redeployment. Higher tax rates can therefore raise firm value![18]

Appendix 3.10 derives and illustrates value and leverage ratios in the two extreme cases. (The leverage ratios comparative statics do not change with the source of the debt benefits τ.)

3.3.4.3 Leverage

We are now ready to proceed to the focus of our paper, corporate leverage, when there are both the traditional financial-distress channel and the novel opportunistic acquisition channel. For what follows, we continue to assume that the benefits of debt (τ) accrue to

---

[18]Note that firms with high debt can pass on the tax shield even when expected earnings are low and they are likely to go bankrupt. However, tax revenue can also improve when reallocational efficiency improves.

shareholders. The debt F_i in our model corresponds to the face value at time 1. Because the expected return on debt is zero in our model, the market value of debt at time 0 is

$$D(F_i) \equiv \begin{cases} F_i & \text{if } F_i < P^* \\ F_i - (1+\phi)\cdot(F_i^2 - \Lambda(F_i)^2)/2 & \text{otherwise .} \end{cases}$$

For a low face value of debt, there is no possibility of default. For a high face value of debt, the expected payout to creditors is equal to the promised payoff, F_i, less the expected loss to creditors.

Below, we will be considering the market value of the low type debt at time 0 in equilibrium, which is denoted $D_L^* \equiv D(F_L^*)$. We will also be considering the industry average market value of the debt at time 0 in equilibrium, which we will call $D_{Ind}^* \equiv h^* \cdot D(F_H^*) + (1-h^*)\cdot D(F_L^*)$.

Absolute Leverage: The left plot in Figure 3.7 shows that industry debt can first increase and then decrease in redeployability η (for low debt benefits τ). For these very low debt benefit values, the financial-distress channel dominates when redeployability is low. At first, when redeployability increases, firms take on more debt. It makes little sense for such firms to speculate on purchasing assets—the assets are simply not valuable enough. Eventually, when redeployability increases further, the potential to buy assets becomes more lucrative, the asset-acquisition channel begins to dominate, and firms again take on less debt. Finally, for higher debt benefits τ, only the asset-acquisition channel matters again. It dominates for all redeployability parameters η. Firms always find it more important to keep leverage low because of the opportunity to pounce on future opportunities.[19]

Implication 4 *For low debt benefits τ and low asset redeployability η, the financial-distress channel dominates. Firms take on more debt when assets become more redeployable. For higher debt benefits τ and higher asset redeployability η, the opportunistic-acquisition channel dominates. Firms take on less debt when assets become more redeployable.*

[19]The plots for the low-debt firm are identical when there is no mixing, and very similar when there is mixing. There is a tiny obscure region in which the low-debt firm may reduce its debt face value when the benefits increase.

Leverage-Value Ratios: There are now two reasons why empirical debt-ratios (the plot on the right in Figure 3.7) may not increase in redeployability. The first effect is the aforementioned endogenous-value effect. Both debt and firm value increase with the direct ongoing debt benefit, and thus the leverage ratio can even decrease in τ. The second effect is the acquisition channel.

Together, a simple linear regression explaining leverage ratios with redeployability proxies is not a powerful test. Instead, a better test would posit a U-shape—first an increasing and then a decreasing effect. When redeployability is low, a small increase in redeployability induces firms to fear distress less and they increase leverage. This is the case regardless of the source of debt gains. When redeployability is high, a small increase in redeployability induces firms to hold out for better acquisition opportunities and they decrease leverage.

A glance at the left and the right plot makes it obvious that the face value of debt and the resulting leverage ratio show completely different behavior. There are wide regions in which the face value of debt increases and the the leverage ratio decreases, and vice-versa.

Implication 5 *Debt face values and leverage ratios can have different comparative statics. One may go up when the other goes down, and vice-versa. This is because parameters effect not only the debt but also the firm value.*

3.3.4.4 Ancillary Implications

Our model can also offer implications on other measures that were not its primary focus. This section provides a sampling.[20]

Credit Spreads: Creditors are indifferent between providing funding and not providing funding to the low type if the credit spread is

$$r(F_L^*) \equiv \frac{F_L^*}{D_L^*} - 1 \,. \tag{3.9}$$

[20] We could also offer further implications on other outcomes (such as on the average values and discounts of assets in production and transfer) that would be more difficult to measure empirically.

The top left plot in Figure 3.8 shows that credit spreads increase when debt benefits are higher. Higher τ encourages firms to take on more debt, which increases the expected loss to creditors. Higher η (redeployability) leads to higher recovery rates and (all else equal) lower credit spreads, but firms may optimally choose higher debt levels which increases the likelihood of default. Although the former effect almost always dominates, resulting in lower spreads when redeployability is greater, there is a very small parameter region in which the credit spread increases when redeployability is greater. Finally (not plotted), just as in Leland (1994), credit spreads may increase or decrease in reorganization costs. Higher ϕ lead to lower recovery rates in the event of default, but also cause firms to choose lower levels of debt, which reduces the likelihood of default.[21]

Asset Liquidation Price (P): All three price-related comparative statics are unambiguous (though they can be quite flat): asset prices increase in redeployability and reorganization costs, and decrease in debt benefits. We already discussed earlier in the context of our model without reorganization costs why the asset price increases with redeployability and decreases with debt benefits. Higher reorganization costs have two competing effects on price: on one hand, they result in greater supply of the asset, because liquidation becomes relatively more desirable than continuing operations in financial distress. On the other hand, they result in greater demand for the asset, because firms take on less debt and therefore have more access to financing. Though not necessarily universal, in our specific model, the latter effect always dominates.

Asset Sales (Q): Asset sales always increase in redeployability and decrease in reorganization costs, but are ambiguous in debt benefits. The dominant effect of greater redeployability is to make the asset more valuable to a potential buyer, resulting in greater demand and higher asset sales. Higher reorganization costs make asset sales more appealing relative to

[21]Our qualitative comparative statics results for credit spreads hold even when we allow creditors to have access to the immediate debt benefits (see Appendix 3.8). Of course, quantitative predictions about credit spreads will depend on whether creditors have access to the immediate debt benefits—which they may in the real world. It is possible to change the model to entertain different assumptions on the disposition of these benefits.

the direct alternative of reorganization. Higher debt benefits increase firm debt. This results both in less demand for the asset (because of tighter financing constraints), and in greater asset supply (because of more firms in trouble). The net effect is ambiguous. Appendix Section 3.9 discusses the cyclicality of asset sales when there is uncertainty in the industry or economy.

Distressed Reorganization Observables: The model also offers secondary predictions for two quantities related to distressed reorganization:

- The liquidation frequency for the low type conditional on being in financial distress is $\Lambda(F_L^*)/F_L^*$ if $F_L^* \geq P$. If $F_L^* < P$, firms will always liquidate and never continue. The plot shows that firms liquidate more often in distress when assets are more redeployable. The dominant effect here is that more redeployable assets have higher liquidation values which makes liquidation more desirable. Higher distress costs reduce continuation values *holding debt fixed*, but higher distress costs also result in lower optimal debt which increases continuation values. The first effect dominates and the conditional liquidation frequency increases in reorganization cost ϕ. Increasing the benefits of debt τ leads to higher debt levels and declining liquidation values. This makes liquidation less attractive and therefore less frequent.

- The expected losses associated with reorganizing the firm, $E_v[\phi \cdot (F_L^* - v) \cdot \mathbb{1}_{\Lambda(F_L^*) \leq v \leq F_L^*}]$ —possibly at least a partial transfer to and thus a partial proxy of the size for the legal reorganization industry—increase in τ, decrease in η, and are ambiguous in ϕ. The dominant effect of increasing τ is to increase debt which increases the likelihood of distress and the dissipative cost of reorganizing and continuing in distress. The dominant effect of increasing η is to increase liquidation values which makes impaired continuation less likely and reduces expected reorganization costs. The effect of ϕ is ambiguous because it increases reorganization costs *holding debt fixed* but reduces optimal debt.

3.4 Discussion and Literature Context

3.4.1 Welfare

Our paper has largely deemphasized welfare implications and government policy prescriptions, because we view the model as too stylized to offer policy prescriptions. Our model assumes production, reallocation, incentive and tax[22] effects as parameters; and we are simply not confident enough to take a stance to what extent these aspects are dissipative or redistributive to some parties elsewhere in the economy.

We are however comfortable to discuss briefly one part of the overall social welfare within the context of our model. This can help to clarify one conceptual aspect of the trade-offs that government should be aware of. How do corporate income taxes—one component of τ—influence re-allocational efficiency?

Figure 3.9 shows that the answer is ambiguous:

Implication 6 *Increases in the benefits of debt—as can be effectuated by tax code changes—can result in socially less or more efficient redeployment activity.*

This is because there is typically an intermediate level of debt, in which asset transfer

[22]To the extent that some of the firm's debt benefits come through the tax shelter (though there are others!), there is a related conceptual puzzle. If, as is widely acknowledged, debt has a potentially negative effect on the stability of firms individually and system-wide, why would the government want to impose them differentially on equity and not on debt? A government could impose taxes on projects instead—for example, in Germany, home ownership is subsidized not through interest deductibility, but through non-recaptured depreciation. We can speculate that default forces more reallocation of resources from less productive to more productive uses; and by increasing the value of debt, the government can calibrate both the equilibrium reallocation frequency and reallocation state dependence. However, debt is a fairly blunt instrument, used by governments that are themselves not great experts about when reallocation is better or worse. The mutual industry-peer externalities discussed in our paper further suggest that it could be a dangerous instrument—if it forces only a few firms to sell, prices are reasonably appropriate, but at some point, feedback effects can reallocate assets less towards the highest-value user of assets and more towards the least-levered users of assets.

activity is socially ideal.[23] Tax policy can then push firms toward or away from this ideal. This is easiest to understand in the context of the total direct debt benefits:

- For low τ, firms choose low leverage, resulting in high demand for liquidated assets. If reorganization costs are high—which makes liquidation more likely in financial distress—this can also result in high supply of assets, and the economy can have too many asset transfers relative to the efficient level. Increasing the tax advantage of debt then pushes firms towards more debt, which helps because it will reduce the expected transfer activity.

- For high τ, firms choose high leverage, resulting in low demand for liquidated assets. The economy has too few asset transfers relative to the efficient level. Increasing the tax advantage of debt further would only push firms towards even more debt and thereby worsen the reallocation.[24]

A reasonable interpretation is that government tax policy towards debt should moderate other debt benefits.

For comparison, in Gale and Gottardi (2015), in which asset sale prices are also endogenous, the thought experiment about the social cost of debt as a tax shelter is different. In their model, in the absence of a corporate debt response (to undo taxes), such taxes would *always* reduce socially beneficial productive operations. Debt, by undoing taxes, tends to increase productive activity and can thereby improve social welfare. Taking the leverage responses of firms into account, the net effect of an increase in taxes on production and thus welfare could be positive or negative. Interestingly, Gale-Gottardi consider a novel policy

[23] Assets are identical and it is always the lowest-use owners who transfer assets to the highest-use owners. Thus, the total quantity transferred is the only metric of relevance.

[24] Similar to this point, when redeployability is low and reorganization costs are high, firms would choose high leverage. This is because maintaining financial flexibility is less valuable when it is unlikely that there will be good buying opportunities later. Again, transfer activity would be too low from a social perspective, and increasing the tax advantage of debt would only hurt more. (The opposite is the case when redeployability is high and reorganization costs are low.)

mechanism—forcing firms to take on more debt. This can in turn induce firms to increase investment voluntarily.

3.4.2 Generalizations of the Model

The most important takeaways of our model are that

1. firms' leverage choices are affected by their peers through the equilibrium price of liquidated assets;
2. indivisibility of assets may result in heterogeneity in leverage strategies;
3. leverage level effects are not isomorphic to leverage-ratio effects;
4. the acquisition channel means that increased asset redeployability can also have a negative effect on leverage, especially when debt benefits and redeployability are high to begin with;
5. and tax policy and non-tax related debt benefits can have ambiguous effects on reallocational efficiency, firm value, and tax receipts. For example, for some large tax rates, a further increase in tax rates can increase both firm value and tax revenues.

To illustrate them, our model had to employ a set of assumptions for tractability, such as the uniform distribution on values; linearity in η, ϕ, and τ; stark integration limits; limited liability and free disposal; limited capital; uncorrelated shocks; no further countervailing important omitted effects (e.g., due to agency or inside information), and so on. None of our takeaways lean especially heavy on specificity in these assumptions, and we would expect the key insights to survive in models in which they are reasonably relaxed. In particular:

- Outside Buyers: Our model is sensitive to the assumption that buying is limited to firms inside the industry. Our qualitative results would continue to hold if there is limited demand from outside the industry—this would increase liquidation values and mitigate, but not eliminate, the incentive to choose lower debt to take advantage of buying opportunities. It would also have a similar effect as an increase in redeployability, η. But if assets are just as valuable outside the industry and potential buyers have practically unlimited capital, then our acquisition channel vanishes, as discussed in Section 3.3.3.3.2. More commonly, neither zero nor infinite capital availability inside the industry, and neither perfect nor useless redeployability outside the industry is likely to be a realistic description; and these unmodeled forces can help push eta towards lower or higher levels.

- Correlated shocks: It could be that all assets in an industry are simultaneously affected by a recession, or that (e.g., consumer taste) shocks help some firms at the same time they hurt others. For example, if shocks are positively correlated, fire sales will be deeper in bad times (more sellers and fewer buyers) and shallower in good times (fewer sellers and more buyers). This may create an incentive to take on less debt initially to take advantage of the great investment opportunities available in bad times, above and beyond the incentive to avoid financial distress oneself. Appendix Section 3.9 sketches an extension of our model to industry uncertainty. It shows that there are parameter regions where reallocation of assets is procyclical and regions where it is counter-cyclical.

- Agency Conflicts: When managers (and equity) have stronger incentives not to declare bankruptcy and even weaker incentives to liquidate (and if creditors cannot renegotiate managers out of collectively inefficient choices, as in Benmelech and Bergman (2008)), then firms would likely be less inclined to liquidate at the same time, given the same amount of debt. However, this would not necessarily be the outcome. In turn, this could have equilibrium repercussions for the optimal level of debt and/or various restrictions written into debt that can enhance the incentives of firms to liquidate. The outcome would likely depend on how extra debt calibrates the relative incentives.

Our paper has endogenous heterogeneity. More realistically, there would be both exogenous heterogeneity and endogenous heterogeneity. We have not modelled differences in behavior across types, however. Firms with higher ex-ante quality could have both more debt capacity and expect to be buyers. It is not clear whether this would lead them to behave differently from lower quality types.

3.4.3 Related Literature

Our model was built around the fundamental tradeoff between taxes and financial-distress costs, first raised in Robichek and Myers (1966). As this encompasses most of the modern theory of corporate capital structure, we can only highlight some work especially close to the assumptions and results of our own paper. Harris and Raviv (1990), Leland (1994), Leland and Toft (1996), Gryglewicz (2011), and many others, have provided the theoretical formalizations to help understand firm tradeoffs and behavior. Industry debt choices have been proposed by Maksimovic and Zechner (1991), Fries, Miller, and Perraudin (1997), and others.[25]

The costs of financial distress were further dissected into components, such as debt overhang (Myers (1977)), the damaged relationships with key stakeholders (Titman (1984)), or reduced market share (Opler and Titman (1994)).

Allen and Gale (1994) and Acharya and Viswanathan (2011) develop models of asset sales in which potential buyers face entry costs or are financially constrained so that equilibrium prices depend on funding availability of industry peers. Unlike our model, these models have specific fixed funding needs, with an endogenous determination of whether they can raise them. (In this sense, they do not choose an optimal capital structure.) Furthermore, assets are divisible in these models; however, we show that when assets are indivisible there may be mixed equilibria in which some firms adopt high-debt strategies to take advantage of tax benefits and others adopt low-debt strategies to take advantage of asset fire sale opportunities.

Gale and Gottardi (2015) offer a theory in which debt is an optimal choice and fire-sale prices are also endogenous.[26] In their model, frictions and especially taxes lead firms

[25] Bolton, Santos, and Scheinkman (2011) motivate preferred purchasers through adverse selection. Asset specificity plays a role in Marquez and Yavuz (2013), though assets have exogenous prices. More specific assets can increase productivity (reducing debt) and increase continuation values (increasing debt).

[26] In Gale and Gottardi (2011), leverage is not a choice that firms consider. (Projects are 100% leverage by assumption.)

to take on too few projects from a social point of view. Debt can reduce the tax burden and thereby enhance the desire of firms to take projects. An endogenous reduction in price upon resale[27] comes into play, because when many firms have taken on too much debt, the induced price reduction then works against this social advantage of debt. As remedy, they propose forcing firms to take on more debt. This induces them to undertake more projects, which in their model is socially valuable. As noted, our model has a different structure, parameters, and focus. It considers social welfare only in passing, because our own model assumes production, reallocation, and tax costs as parameters, and we are less confident about the dissipative/redistributive cost-benefit issues for them.

A number of empirical papers have provided evidence about the existence and nature of these fire sales. Asquith, Gertner, and Scharfstein (1994) showed that financially-distressed firms often liquidate assets at discounts to fundamental value. Pulvino (1998) showed that there are periods in which many airlines were hit by negative shocks at the same time, how this depressed airplane prices, and how financially unconstrained airlines then increased their buying activity, while constrained airlines did not. Acharya, Bharath, and Srinivasan (2007) investigated this effect more generally. Taking this yet a step further, Benmelech, Garmaise, and Moskowitz (2005) showed that firms take on more debt when assets are easier to redeploy. Rajan and Ramcharan (2016) show how financial intermediation failures have reduced land sale values through fire sales.

They interpreted their findings as support for an optimal capital structure theory, in which assets that were more redeployable allowed industries to take on more debt.

3.4.4 Relation to Shleifer-Vishny 1992

Like Jensen and Meckling (1976) did for agency issues, the seminal Shleifer and Vishny (1992) paper laid out a research agenda for corporate behavior when asset prices could be depressed by the need of many firms to sell simultaneously and the resulting fire-sale

[27]Assets are as productive to buyers as they were to sellers. Sales are costly to the firm, but not to the economy.

liquidations.[28] However, Shleifer-Vishny do not present an economic model in the traditional sense. Their paper model lays out a set of inequalities that an equilibrium should satisfy in which two firms' debt choices could influence one another. It offers neither a specific solution to these conditions nor any comparative statics. This model sketch is then followed by an insightful discussion of economic possibilities. However, much of this discussion does *not* follow from their model. It is therefore not clear from their paper whether it is possible to construct a model of firms that exhibit the kinds of behavior they conjecture. Our own paper has provided such a model.

Their Section II describes a three-period model with two industry firms plus one outsider. Both firms choose the minimum debt levels that eliminate overinvesting in prosperous times. The model describes a set of constraints[29] under which each firm is either a potential purchaser for the other firm or a bystander, suffering from its cash constrained inability to buy the selling firm's assets. Other model parameters can push this within-industry value above or below the external asset value. The price of the liquidating asset is the lower of the two. Both firms can calculate the other-firm price in advance, which they can take into consideration when they choose whether to have debt (and avoid the agency costs) or not to have debt (and avoid costly liquidation).

With only one outsider and one insider competing, the opportunistic buying decision is limited to whether each firm wants to have low enough debt to beat the external price (be the acquiror) if the other firm runs into difficulties, or not. Any buying firm always pays the

[28]Duffie (2010) went even further, attributing temporarily depressed prices not just to firms and industry assets, but even to financial claims in wide distribution.

[29]The model has 12 exogenous firm parameters (4 cash flow parameters and 2 investment flow parameters, per firm); 1 internal and 1 external asset value parameter; 1 probability that governs the state of the economy; 4 endogenous debt parameters (short-term and long-term, per firm), resulting in the *key* resulting maximum value that a buyer can pay, given own debt overhang; and 14 equality and inequality constraints (guaranteeing such conditions as debt overhang being optimal, firms needing to raise capital, control of agency in good times being more important than liquidation in bad times, etc.). 5 conditions are redundant. There are also 7 other explicit assumptions and some implicit assumptions (like $d \geq 0$). It is not difficult to find 15 parameters that satisfy their 14 conditions.

outside value. The other firm's high or low debt choice thus has a (binary) influence on the value of the firm in distress. If there is any competition in the external market, the price of the liquidating firm would become completely independent of industry debt choices (i.e., the other firm). Instead, the price would always be determined by the external market value alone. To resurrect feedback from the industry price to debt choices back into the model would require the introduction of multiple insiders, who would compete against one another when setting the price. The dynamics of such a model would be more complex and are not at all obvious.

The S-V model does not derive comparative statics from model solutions. In contrast, our model derives specific optimal debt levels, equilibrium prices, and values for firms in industries that are subject to distress and tax costs. It thus offers comparative statics with respect to taxes and distress costs, and especially with respect to redeployability. This yielded our key result about where *debt can go up or down with redeployability*, in contrast to Williamson (1988) and Harris and Raviv (1990). The "important general principle: optimal leverage or debt capacity falls as liquidation value falls" (p.1354) in Shleifer-Vishny holds only for the seller's leverage when a threshold is crossed and he chooses discontinuously to take on no debt. Our model characterized how leverage choices may change continuously in reaction to liquidation prices due to both buying and selling concerns.

In the region of the S-V model where two equilibria exist, either the buyer chooses a low (zero) debt level and the seller a high one, or vice versa. The heterogeneity is exogenous—the model cannot accommodate identical firms. In this region "it would not be an equilibrium either for both firms to have a lot of debt or for neither to have debt. This case of two equilibria suggests the notion of an industry debt capacity..." (p.1354), similar to Miller (1977). Of course, firms are not exactly homogeneous ex-ante, as characterized by our model. However, our model can explain (a) how the notion of an industry debt level matters even when all firms choose the same debt and (b) when otherwise similar firms may or may not end up choosing different debt strategies. The latter choice depends crucially on the indivisibility of the asset. With divisible assets, all identical firms would choose the same debt.

3.4.5 Summary

Table 3.3 highlights the key difference between our model and the most closely related papers in the literature. Primarily,

1. Our model has offered a negative comparative static with respect to redeployability.
2. Our model has generated endogenous heterogeneity among firms that are ex-ante homogeneous, with an important link to asset indivisibility.

Secondarily,

3. Our model has offered comparative statics on leverage (D/V), not just on debt levels.
4. It is among very few models in which prices are endogenous.

3.5 Conclusion

Our paper has sketched a model in which firms could anticipate and participate in industry asset sales, with more levered firms as sellers and less levered firms as buyers. This turns prices into mediators of industry leverage interactions, and ambiguates the role of asset redeployability. When redeployability is low, an increase therein induces firms to take on more debt in order to take advantage of higher fire sales prices as potential sellers—as in the earlier literature. However, when redeployability is high, an increase therein induces firms to take on less debt in order to take advantage of fire sales as potential buyers.

3.6 Appendix: Proof of Theorem 1

Each firm is competitive and takes the price, P, as given. The marginal benefit of debt is τ for all values of debt, F. The marginal cost of debt falls into three regions:

1. For $F \in [0, P(1/\eta - 1)]$ the marginal cost is zero

2. For $F \in [P(1/\eta - 1), 1 - P]$ the marginal cost is $\eta(F + P) - P$

3. For $F \in [1 - P, 1]$ the marginal cost is zero

In Region 1, increasing debt is not costly to the firm because the marginal project is negative NPV. In Region 2, increasing debt is costly because the firm must forego positive NPV projects. In Region 3, the debt level is so high that the firm cannot finance the acquisition of the asset even if it is the highest productivity type, $V_i = 1$. Therefore, increasing debt further results in no additional costs to the firm.

Since the marginal benefit of debt is positive (equal to τ) it follows that it is never optimal for the firm to choose a debt level in Region 1, i.e., $F \leq P/\eta - P$. Furthermore, since the marginal cost of debt jumps down to zero at $F = 1 - P$ there may be a mixed equilibrium in which some firms choose debt in Region 2 while others choose debt in Region 3. Clearly, in such an equilibrium, firms in Region 3 will choose $F_H = 1$ since the marginal benefit of debt exceeds the marginal cost for all debt choices in Region 3.

The first-order condition for an optimal (interior) debt choice is:

$$F_L(P) = P/\eta - P + \tau/\eta.$$

The second-order condition is clearly satisfied.

3.6.1 Pure-strategy interior equilibrium

Without reorganization costs the firm liquidates at time 1 if and only if $V_i \leq P$, therefore, the supply of the asset is P. In a pure-strategy equilibrium, the demand is $1 - P - F_L(P)$, therefore, we have the unique market clearing price

$$P^* = \frac{\eta - \tau}{1 + \eta}$$

and the unique interior debt choice

$$F_L^* = \frac{2\tau + 1 - \eta}{1 + \eta}.$$

3.6.2 Mixed-strategy equilibrium

We now consider the possibility of a mixed equilibrium. The value of a firm choosing $F_L(P)$ is

$$\begin{aligned} V_L(P) &= 0.5 \cdot (1 + P^2) + \int_{P+F_L}^{1} (\eta V - P) dV + \tau F_L \\ &= 0.5 \cdot (1 + P^2) + \frac{(\eta^2 + P^2 - 2P\eta - \tau^2)}{2\eta} + \tau F_L, \end{aligned}$$

and the value of a firm choosing $F_H = 1$ is

$$V_1(P) = 0.5 \cdot (1 + P^2) + \tau.$$

Setting $V_L(P) = V_1(P)$ and solving for P yields the unique market clearing price in a mixed-strategy equilibrium:

$$P = \eta - \tau + \eta\tau - \sqrt{\eta^2\tau^2 + 2\eta\tau(\eta - \tau)}.$$

There is another candidate P, but it is greater than the pure-strategy equilibrium price. We know that this price could never be supported in equilibrium, as introducing high-debt firms both reduces the demand and increases the supply of the liquidated asset.

We now find the boundaries of the mixed-strategy equilibrium. Since prices are continuous we know $P^*(\tau^c) = P(\tau^c)$ (i.e., prices in the pure-strategy and mixed-strategy equilibria are equal at the boundaries). Solving for τ^c gives:

$$\tau^c = \frac{2\eta^2 + \eta \pm \eta(1 + \eta)}{3\eta + 2} \Rightarrow \{\tau_{c_1}, \tau_{c_2}\} = \left\{ \frac{\eta^2}{3\eta + 2}, \eta \right\}.$$

Therefore, for $\tau \in \left(\frac{\eta^2}{3\eta+2}, \eta \right]$ there is a mixed-strategy equilibrium with proportion h^* of firms choosing $F_H = 1$, proportion $1 - h^*$ of firms choosing $F_L(P^*) = P^*/\eta - P^* + \tau/\eta$, and the price

$$P^* = \eta - \tau + \eta\tau - \sqrt{\eta^2\tau^2 + 2\eta\tau(\eta - \tau)}.$$

The supply of the asset is P^* and in a mixed-strategy equilibrium the demand is $(1 - h^*) \cdot (1 - P^* - F_L)$, therefore, we can solve for the unique proportion:

$$h^* = \frac{1 - 2P^* - F_L(P^*)}{1 - P^* - F_L(P^*)}.$$

3.6.3 Pure-strategy extreme equilibrium

For $\tau > \eta$ there is a pure-strategy equilibrium in which all firms choose $F^* = 1$ and the equilibrium price is $P^* = 0$. In this region, the marginal benefit always exceed the marginal cost. In this region, the demand for the asset is zero, as the financing constraint always binds. The supply is also zero, since the manager always prefers to keep the asset worth $0 \leq V_i$ instead of selling it for nothing.

3.6.4 Uniqueness

All together, we can characterize which type of symmetric equilibrium will obtain by looking at the exogenous parameters. For $\tau \in \left[0, \frac{\eta^2}{3\eta+2}\right]$, we have a unique, pure-strategy interior equilibrium. For $\tau \in \left(\frac{\eta^2}{3\eta+2}, \eta\right]$, we have a unique, mixed-strategy equilibrium. For $\tau \in (\eta, 1]$, we have a unique, pure-strategy extreme equilibrium. It cannot be that there is a mixed-strategy equilibrium in either of the pure-strategy regions, as it would require $h^* < 0$ or $h^* > 1$ to support the equilibrium price, which is not possible. Also, there cannot be either of the pure-strategy equilibria in the mixed-strategy region. Here, the fraction h^* is chosen to make firms indifferent between high and low debt. If all the firms chose high debt, it would cause prices to fall and make F_L more attractive. Conversely, if all of the agents selected low debt, prices would rise and F_H would be preferable.

We also established that the equilibrium debt levels and prices are unique functions of the exogenous parameters. Therefore, for any given set of exogenous parameters, we can identify the unique symmetric equilibrium.

3.7 Appendix: Extension of Theorem 2 to the Full Parameter Space

3.7.1 Statement

In the text, we covered a limited parameter space for illustration. This appendix states the theorem and proof for the model's complete parameter space.

$$\begin{aligned}
\text{Let } \tau_0 &= \frac{\phi(1-2\eta)}{(1+\eta)(1+\phi)+\phi(1-2\eta)}, \\
\tau_1 &= \frac{2\eta^2(\phi+1)-\phi+\eta(\phi+1)-\sqrt{(\eta+1)^2(\phi+1)\left(\eta^2\phi+\eta^2-2\phi-\eta\phi\right)}}{3\eta(\phi+1)+3\phi+2}, \\
\tau_2 &= \frac{(2\eta-1)(2\eta+\phi+2\eta\phi)}{2+3\phi} \\
&- \frac{\sqrt{(2\eta-1)^2(1+\phi)(4\eta^2(1+\phi)+\eta\phi-2(\eta+\phi))}}{2+3\phi}, \\
\tau_3 &= \frac{2\eta^2(1-\phi^2)+\eta(1+9\phi+\phi^2-5\phi^3)+2\phi+12\phi^2+7\phi^3}{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2} \\
&- \frac{\sqrt{(1+\phi)(1+\eta+5\phi-\eta\phi)^2(\eta^2(1+\phi)(1+3\phi^2)-3\eta\phi^2(2+\phi)-2\phi^3)}}{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2}, \\
\tau_4 &= \frac{2\eta^2(1-\phi^2)+\eta(1+9\phi+\phi^2-5\phi^3)+2\phi+12\phi^2+7\phi^3}{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2} \\
&+ \frac{\sqrt{(1+\phi)(1+\eta+5\phi-\eta\phi)^2(\eta^2(1+\phi)(1+3\phi^2)-3\eta\phi^2(2+\phi)-2\phi^3)}}{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2}.
\end{aligned}$$

Region 1: $\eta \geq 1/2$ **and** $\phi < \frac{3\eta-2}{6-3\eta}$

- If $0 \leq \tau \leq \tau_1$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta-\tau}{1+\eta}$ in which all firms choose $F_L^* = \frac{1-\eta+2\tau}{1+\eta}$.

- If $\tau_1 < \tau \leq \tau_2$ there exists a unique, mixed-strategy equilibrium with price

$$\begin{aligned}
P^* &= \frac{\phi\eta-(1+\phi)[\tau-\eta(1+\tau)]}{1+\phi(1+\eta)} \\
&- \frac{\sqrt{\eta(\phi+1)(2\tau(\eta\phi+(\eta-\tau)(1+\phi))+\eta\tau^2(\phi+1)-\phi(1+\tau-\eta)^2)}}{1+\phi(1+\eta)}.
\end{aligned}$$

in which the proportion $1-h^*$ of firms choose $F_L^* = \frac{\tau}{\eta} + \frac{(1-\eta)}{\eta}P^*$ and the proportion h^* of firms choose $F_H^* = 1$, where

$$h^* = \frac{(1+\phi)\cdot(\eta-\tau-(1+\eta)\cdot P)}{\eta\phi+(1+\phi)(\eta-\tau)-(1+\phi(1+\eta))\cdot P}.$$

- If $\tau_2 < \tau \leq \tau_4$ there exists a unique, mixed-strategy equilibrium with price

$$P^* = \frac{\phi(1+2\phi(1-\tau)-3\tau)+\eta(1+\phi)(1+\tau+\phi(2+\tau))-\tau}{1+(6-3\eta)\phi(1+\phi)} - \frac{\sqrt{(1+\phi)(\eta+\phi+\eta\phi)\begin{pmatrix}3\eta^2\phi(1+\phi)-2(\phi(\tau-1)+\tau)^2\\ +\eta(\phi(\tau-1)+\tau)(2+\phi(\tau-1)+\tau)\end{pmatrix}}}{1+(6-3\eta)\phi(1+\phi)}. \tag{3.10}$$

in which the proportion $1-h^*$ of firms choose $F_L^* = \frac{\tau(1+\phi)}{(\eta+\phi+\eta\phi)} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P^*$ and the proportion h^* of firms choose $F_H^* = 1$, where

$$h^* = \frac{(1+\phi)\cdot(\eta+\phi+\eta\phi-\tau(1+2\phi)-P\cdot(1+\eta+5\phi-\eta\phi))}{(1+2\phi)\cdot(\eta+\phi+\eta\phi-\tau(1+\phi))-P\cdot(1+5\phi-\eta\phi+5\phi^2-\eta\phi^2)}. \tag{3.11}$$

- If $\tau_4 < \tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$ in which all firms choose $F_L^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$.

Region 2: $\eta \geq 1/2$**,** $\phi \geq \frac{3\eta-2}{6-3\eta}$**, and** $\eta^2(1+\phi)(1+3\phi^2)-3\eta\phi^2(2+\phi)-2\phi^3 \geq 0$

- If $0 \leq \tau \leq \frac{2\eta-1}{3}$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta-\tau}{1+\eta}$ in which all firms choose $F_L^* = \frac{1-\eta+2\tau}{1+\eta}$.

- If $\frac{2\eta-1}{3} < \tau \leq \tau_3$ or if $\tau_4 < \tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$ in which all firms choose $F_L^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$.

- If $\tau_3 < \tau \leq \tau_4$ there exists a unique, mixed-strategy equilibrium with price (3.10) in which the proportion $1-h^*$ of firms choose $F_L^* = \frac{\tau(1+\phi)}{(\eta+\phi+\eta\phi)} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P^*$ and the proportion h^* of firms choose $F_H^* = 1$, where h^* is (3.11).

Region 3: $\eta \geq 1/2$**,** $\phi \geq \frac{3\eta-2}{6-3\eta}$ **and** $\eta^2(1+\phi)(1+3\phi^2)-3\eta\phi^2(2+\phi)-2\phi^3 < 0$

- If $0 \leq \tau \leq \frac{2\eta-1}{3}$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta-\tau}{1+\eta}$ in which all firms choose $F_L^* = \frac{1-\eta+2\tau}{1+\eta}$.

- If $\frac{2\eta-1}{3} < \tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$ in which all firms choose $F_L^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$.

Region 4: $\eta < 1/2$ **and** $\eta^2(1+\phi)(1+3\phi^2) - 3\eta\phi^2(2+\phi) - 2\phi^3 \geq 0$

- If $0 < \tau \leq \tau_0$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta(1-\tau)}{1+\eta}$ in which all firms choose $F_L^* = \frac{\tau(1+\eta+\phi)+\phi\eta}{\phi(1+\eta)}$.

- If $\tau_0 < \tau \leq \tau_3$ or if $\tau_4 \leq \tau \leq \frac{\phi+\eta(1+\phi)}{1+2\phi}$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$ in which all firms choose $F_L^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$.

- If $\tau_3 < \tau \leq \tau_4$ there exists a unique, mixed-strategy equilibrium with price (3.10) in which the proportion $1-h^*$ of firms choose $F_L^* = \frac{\tau(1+\phi)}{(\eta+\phi+\eta\phi)} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P^*$ and the proportion h^* of firms choose $F_H^* = 1$, where h^* is (3.11).

Region 5: $\eta < 1/2$ **and** $\eta^2(1+\phi)(1+3\phi^2) - 3\eta\phi^2(2+\phi) - 2\phi^3 < 0$

- If $0 < \tau \leq \tau_0$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta(1-\tau)}{1+\eta}$ in which all firms choose $F_L^* = \frac{\tau(1+\eta+\phi)+\phi\eta}{\phi(1+\eta)}$.

- If $\tau_0 < \tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$ there exists a unique, pure-strategy equilibrium with price $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$ in which all firms choose $F_L^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$.

Region 6:

- If $\frac{\eta+\phi+\eta\phi}{1+2\phi} < \tau \leq 1$ there exists a unique, pure-strategy equilibrium with price $P^* = 0$ in which all firms choose $F_L^* = \min\left\{1, \frac{\tau(1+\phi)}{\eta(1+\phi)+\phi}\right\}$.

3.7.2 Proof

Proof when $\eta \geq 1/2$

We first consider the case $\eta \geq 1/2$. Each firm is competitive and takes the price, P, as given. For now, we assume $P > 0$. We will consider the possibility that $P = 0$ later in the proof. The marginal benefit of debt is τ for all values of debt, F. If $P > 0$ the marginal cost of debt falls into four regions:

1. For $F \in [0, P \cdot (1/\eta - 1)]$ the marginal cost of debt is 0
2. For $F \in (P \cdot (1/\eta - 1), P]$ the marginal cost of debt is $\eta F - P \cdot (1 - \eta)$
3. For $F \in (P, 1 - P)$ the marginal cost of debt is $\eta F - P \cdot (1 - \eta) + (F - P) \cdot \phi/(1 + \phi)$
4. For $F \in [1 - P, 1]$ the marginal cost of debt is $(F - P) \cdot \phi/(1 + \phi)$

Importantly, the marginal cost of debt is weakly increasing and continuous in F over the first three regions but then jumps down at $F = 1 - P$ (since the financing constraint is no longer binding) after which it increases again. Therefore, for any given marginal benefit τ, there are at most two possible optimal debt choices, one where $F < 1-P$, and one where $1-P < F \leq 1$. Consequently, there is the possibility of both pure-strategy equilibria and mixed-strategy equilibria in which some firms choose low debt and others choose high debt. We must consider three cases: (i) $0 \leq \tau < (2\eta - 1) \cdot P$, (ii) $(2\eta-1) \cdot P \leq \tau \leq \eta - P + (1-2P)\phi/(1+\phi)$, and (iii) $\eta - P + (1 - 2P)\phi/(1 + \phi) < \tau \leq 1$.

Case 1: $0 \leq \tau < (2\eta - 1) \cdot P$

If $\tau < (2\eta - 1) \cdot P$ then firms choose either $F_L \in [P \cdot (1/\eta - 1), P)$ or $F_H \in [1 - P, 1]$ where $F_L = \tau/\eta + P \cdot (1/\eta - 1)$ and $F_H = \min\{1, P + \tau \cdot (1 + \phi)/\phi\}$.

Pure-strategy equilibria

There cannot exist a symmetric equilibrium with $P > 0$ in which all firms choose F_H because the aggregate demand for the risky asset would be zero but the supply is positive $[(P + \phi F)/(1 + \phi)]$. Therefore, if there exists a symmetric equilibrium in the case $\tau < (2\eta - 1) \cdot P$, then all firms choose $F_L = \tau/\eta + P \cdot (1/\eta - 1)$. The demand for the liquidated asset is then $1 - P - F_L$ and the supply of the liquidated asset is P since for $F_L < P$ which implies $\tau < (2\eta - 1) \cdot P$ it is optimal to liquidate the asset for all $V \leq P$. Equating supply and demand gives the unique market clearing price $P^* = (\eta - \tau)/(1 + \eta)$.

Importantly, note that if $P^* = (\eta - \tau)/(1 + \eta)$ then we must have $\tau < (2\eta - 1)/3$ to be in a symmetric equilibrium in the case $\tau < (2\eta - 1) \cdot P$.

Mixed-strategy equilibria

There is also the possibility of a mixed-strategy equilibrium (a fraction of firms choosing F_L and the remaining fraction of firms choosing F_H). Firms choosing F_L have ex ante value

$$V_L = \int_0^P P dV + \int_P^1 V dV + \int_{P+F_L}^1 (\eta V - P) dV + \tau \cdot F_L$$

and by substituting $F_L = \tau/\eta + P \cdot (1/\eta - 1)$ yields

$$V_L = 0.5(\eta + (P - 1)^2 + (P + \tau)^2/\eta - 2P\tau).$$

Firms choosing F_H have ex ante value

$$V_H = \int_0^\Lambda P dV + \int_\Lambda^{F_H} [V - \phi \cdot (F_H - V)] dV + \int_{F_H}^1 V dV + \tau \cdot F_H.$$

where $\Lambda = (P + \phi F)/(1 + \phi)$.

If $F_H = P + \tau \cdot (1 + \phi)/\phi$ then ex ante value is

$$V_H = 0.5 \cdot (P + \tau)^2 + 0.5 \cdot \tau^2/\phi + 0.5,$$

but if $F_H = 1$ then ex ante value is

$$V_1 = 0.5 \cdot (P + \phi)^2/(1 + \phi) + 0.5 \cdot (1 - \phi) + \tau.$$

The following result shows that in a mixed-strategy equilibrium the high-type always chooses $F_H = 1$.

Lemma 1: If $\tau < (2\eta - 1)/3$ then in a mixed-strategy equilibrium the high type chooses $F_H = 1$.

Proof: Proof by contradiction. Suppose $F_H = P + \tau \cdot (1+\phi)/\phi < 1$. Then we have:

$$G(P) \equiv V_L(P) - V_H(P) = 0.5 \cdot (\eta + (P-1)^2 + (P+\tau)^2 \cdot (1-\eta)/\eta - 2P\tau - \tau^2/\phi - 1).$$

Note, $G'(P) = P/\eta - 1 + \tau \cdot (1/\eta - 2) \leq P/\eta - 1 \leq 0$ where the first inequality follows from our assumption that $1/\eta - 1 \leq 1$ and the second from the fact that $P \leq \eta$ as the price for the asset will never exceed its maximum value. Therefore, G(P) is decreasing in P. Furthermore, in a mixed-strategy equilibrium the price is bounded above by the pure-strategy equilibrium price (i.e. $P \leq (\eta - \tau)/(1+\eta)$) because the introduction of some high-debt firms both reduces the demand and increases the supply of the liquidated asset. Therefore,

$$\begin{aligned}
G(P) &\geq G((\eta-\tau)/(1+\eta)) \\
&= 0.5 \cdot \left[\left(\frac{1+\tau}{1+\eta}\right)^2 (1+\eta-\eta^2) + \eta - 2\tau\left(\frac{\eta-\tau}{\eta+1}\right) - \frac{\tau^2}{\phi} - 1\right] \\
&> 0.5 \cdot \left[\left(\frac{1+\tau}{1+\eta}\right)^2 (1+\eta-\eta^2) + \eta - 2\tau\left(\frac{\eta-\tau}{\eta+1}\right) - \tau(1-\tau) - 1\right] \\
&\geq 0 \quad \forall \tau
\end{aligned}$$

where the second inequality follows from the fact that if $F_H^* < 1$ and $P > 0$ then $\tau < \phi/(1+\phi)$ which implies $\phi > \tau/(1-\tau)$; and the third inequality is easily verified numerically. But this contradicts the optimality of F_H^*.

By Lemma 1, we must only compare V_L to V_1 to find a mixed-strategy equilibrium.

Conjecture the existence of a pure-strategy equilibrium in which $F_L = \tau/\eta + P \cdot (1/\eta - 1)$ and $P^* = (\eta - \tau)/(1+\eta)$. Substituting P^* into our expressions for V_L and V_1 implies:

$$H(\tau) \equiv V_L(\tau) - V_1(\tau) = \frac{\eta(\eta-\tau)^2(\phi+1) + \phi(\tau+1)^2 - 2(\eta+1)(\phi+1)\tau(\eta-\tau)}{2(\eta+1)^2(\phi+1)}$$

Therefore,

$$H'(\tau) = \frac{\phi - \eta(1+\phi) - 2\eta^2(1+\phi) + (2+3\phi+3\eta(1+\phi))\tau}{(1+\eta)^2(1+\phi)}$$

and

$$H''(\tau) = \frac{2 + 3\phi + 3\eta(1+\phi)}{(1+\eta)^2(1+\phi)}$$

Note that $H'\left(\frac{2\eta-1}{3}\right) = -\frac{2}{3(1+\eta)(1+\phi)} < 0$ and $H''(\tau) > 0$ for all τ. Therefore, $H'(\tau) < 0$ for all $\tau \leq \frac{2\eta-1}{3}$.

Also, $H\left(\frac{2\eta-1}{3}\right) = \frac{2-3\eta+\phi\cdot(6-3\eta)}{18(1+\phi)} \geq 0$ if and only if $\phi \geq \frac{3\eta-2}{6-3\eta}$. Note that if $\eta < \frac{2}{3}$, then $\frac{3\eta-2}{6-3\eta} < 0$ and $H\left(\frac{2\eta-1}{3}\right) \geq 0$ for any ϕ.

Therefore, if $\phi \geq \frac{3\eta-2}{6-3\eta}$ and $\tau \leq \frac{2\eta-1}{3}$ then $H(\tau) \geq 0$ and all firms optimally choose $F_L = \tau/\eta + P\cdot(1/\eta - 1)$ and the conjectured equilibrium price of $P = (\eta-\tau)/(1+\eta)$ is confirmed by the firms' debt decisions.

However, if $\phi < \frac{3\eta-2}{6-3\eta}$, then the conjectured pure-strategy equilibrium is confirmed only for $\tau \leq \tau_1$ where $H(\tau_1) \equiv 0$. For $\tau > \tau_1$ there is a mixed-strategy equilibrium in which some firms choose $F_L = \tau/\eta + P\cdot(1/\eta - 1)$ and others choose $F_H = 1$ (by Lemma 1). Solving $F(\tau_1) = 0$ yields

$$\tau_1 = \frac{2\eta^2(\phi+1) - \phi + \eta(\phi+1) - \sqrt{(\eta+1)^2(\phi+1)\left(\eta^2\phi + \eta^2 - 2\phi - \eta\phi\right)}}{3\eta(\phi+1) + 3\phi + 2}$$

(Note: There is another solution to $H(\tau_1) = 0$ where $H(\tau)$ again becomes positive beyond that point. However, $H'(\tau) < 0$ for all $\tau \leq \frac{2\eta-1}{3}$ so we know $H(\tau) < 0$ for all $\tau_1 < \tau < \frac{2\eta-1}{3}$.)

For $\tau > \tau_1$ there is a unique, mixed-strategy equilibrium which is constructed by finding P that equates $V_L = V_1$, which is quadratic in P. There are two solutions, but only one where P is less than the pure-strategy price (which must be true in equilibrium as argued above) and it is

$$P^* = \frac{\phi\eta - (1+\phi)[\tau - \eta(1+\tau)]}{1+\phi(1+\eta)} - \frac{\sqrt{\eta(\phi+1)(2\tau(\eta\phi + (\eta-\tau)(1+\phi)) + \eta\tau^2(\phi+1) - \phi(1+\tau-\eta)^2)}}{1+\phi(1+\eta)},$$

Let h be the fraction of firms choosing $F = 1$. The demand for the risky asset is then $(1-h)\cdot(1-P-F_L)$ and the supply of the risky asset is $(1-h)\cdot P + h\cdot(P+\phi\cdot 1)/(1+\phi)$, therefore, market clearing requires a unique proportion of high debt firms:

$$h^* = \frac{(1+\phi)\cdot(\eta-\tau-(1+\eta)\cdot P)}{\eta\phi+(1+\phi)(\eta-\tau)-(1+\phi(1+\eta))\cdot P},$$

Finally, if there is a mixed-strategy equilibrium at the upper boundary we know that $P < (\eta-\tau)/(1+\eta)$ and therefore $\tau < (2\eta-1)/3$ at the boundary. Equating $\tau_2 = (2\eta-1)\cdot P^*(\tau_2)$ yields the upper boundary in this case:

$$\begin{aligned}\tau_2 &= \frac{(2\eta-1)(2\eta+\phi+2\eta\phi)}{2+3\phi}\\ &- \frac{\sqrt{(2\eta-1)^2(1+\phi)(4\eta^2(1+\phi)+\eta\phi-2(\eta+\phi))}}{2+3\phi}\end{aligned}$$

(Note: There is another root but it is greater than $(2\eta-1)/3$ when $\eta \geq 1/2$ so we can ignore it.)

Case 2: $(2\eta-1)\cdot P \leq \tau \leq \eta - P + (1-2P)\phi/(1+\phi)$

If $(2\eta-1)\cdot P \leq \tau \leq \eta - P + (1-2P)\phi/(1+\phi)$ then firms choose either $F_L \in [P, 1-P)$ or $F_H \in [1-P, 1]$ where $F_L = \frac{\tau(1+\phi)}{\eta+\phi+\eta\phi} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P$ and $F_H = \min\{1, P+\tau\cdot(1+\phi)/\phi\}$.

Pure-strategy equilibria

Again, there cannot exist a pure-strategy equilibrium with $P > 0$ in which all firms choose F_H because the aggregate demand for the risky asset would be zero but the supply is positive $[(P+\phi F)/(1+\phi)]$. Therefore, in a pure-strategy equilibrium firms choose

$$F = \frac{\tau(1+\phi)}{\eta+\phi+\eta\phi} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P.$$

The demand for the liquidated asset is $1-P-F$ and the supply of the liquidated asset is

$$\Lambda = \frac{P+\phi F}{1+\phi} = \frac{P\cdot((1-\phi)\eta+2\phi)+\phi\tau}{\eta+\phi+\eta\phi}.$$

Equating supply and demand gives the unique market clearing price

$$P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}.$$

Substituting into the expression for F yields

$$F^* = \frac{1 - \eta - \eta\phi + 2\phi + \tau(2 + \phi)}{1 + \eta + 5\phi - \eta\phi}.$$

Mixed-strategy equilibria

As in Case 1, there is the possibility of a mixed-strategy equilibrium. Firms choosing F_L have ex ante value

$$V_L = \int_0^\Lambda P dV + \int_\Lambda^{F_L} [V - \phi(F_L - V)] dV + \int_{F_L}^1 V dV + \int_{P+F_L}^1 (\eta V - P) dV + \tau \cdot F_L$$

and by substituting $F_L = \frac{\tau(1+\phi)}{\eta+\phi+\eta\phi} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)} P$ yields

$$\begin{aligned} V_L &= \left(\frac{1 + \eta + 6\phi - 3\eta\phi}{2(\eta + \phi + \eta\phi)} \right) P^2 - \left(1 - \frac{\tau(1 - \eta + 2\phi - \eta\phi)}{\eta + \phi + \eta\phi} \right) P \\ &+ \left(1 + \frac{\eta(\eta - 1) + \phi(\eta^2 - 1) + \tau^2(1 + \phi)}{2(\eta + \phi + \eta\phi)} \right). \end{aligned} \tag{3.12}$$

Firms choosing F_H have ex ante value

$$V_H = \int_0^\Lambda P dV + \int_\Lambda^{F_H} [V - \phi \cdot (F_H - V)] dV + \int_{F_H}^1 V dV + \tau \cdot F_H.$$

If $F_H = P + \tau \cdot (1 + \phi)/\phi$ then ex ante value is

$$V_H = 0.5 \cdot (P + \tau)^2 + 0.5 \cdot \tau^2/\phi + 0.5 \tag{3.13}$$

but if $F_H = 1$ then ex ante value is

$$V_1 = 0.5 \cdot (P + \phi)^2/(1 + \phi) + 0.5 \cdot (1 - \phi) + \tau. \tag{3.14}$$

The next result shows that in a mixed-strategy equilibrium the high type always chooses $F_H = 1$.

Lemma 2: In a mixed-strategy equilibrium the high type chooses $F_H = 1$.

Proof: Proof by contradiction. Suppose $F_H^* = P + \tau \cdot (1 + \phi)/\phi < 1$. Since we assume $P > 0$, this implies $\tau < \phi/(1 + \phi)$. We have:

$$\begin{aligned} G(P) &\equiv V_L(P) - V_H(P) \\ &= \left(\frac{1 + (5 - 4\eta)\phi}{2(\eta + \phi + \eta\phi)} \right) P^2 - \left(1 - \frac{\tau(1 + \phi)(1 - 2\eta)}{\eta + \phi + \eta\phi} \right) P + \left(\frac{\eta}{2} - \frac{\eta\tau^2(1 + \phi)^2}{2\phi(\eta + \phi + \eta\phi)} \right) \end{aligned}$$

Therefore,

$$\begin{aligned} G'(P) &= \left(\frac{1+(5-4\eta)\phi}{\eta+\phi+\eta\phi}\right)P - \left(1-\frac{\tau(1+\phi)(1-2\eta)}{\eta+\phi+\eta\phi}\right) \\ &\leq \left(\frac{1+(5-4\eta)\phi}{\eta+\phi+\eta\phi}\right)\left(\frac{\eta(1+\phi)+\phi-\tau(1+2\phi)}{(5-\eta)\phi+1+\eta}\right) - \left(1-\frac{\tau(1+\phi)(1-2\eta)}{\eta+\phi+\eta\phi}\right) \\ &\leq 0 \end{aligned}$$

where the first inequality follows from the fact that in a mixed-strategy equilibrium the equilibrium price is less than the pure-strategy equilibrium price (i.e., $P \leq \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$) and the second inequality holds for all $\eta \geq 1/2$. Therefore,

$$\begin{aligned} G(P) &\geq G\left(\frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}\right) \\ &> 0 \quad \forall \tau < \phi/(1+\phi). \end{aligned}$$

where the last inequality is easily verified numerically. But this contradicts the optimality of F_H^*.

By Lemma 2, we must only compare V_L to V_1 to find a mixed-strategy equilibrium. We have

$$\begin{aligned} V_L - V_1 &= \frac{(\phi(\phi+1)(2-3\eta)+(2\phi+1)^2)}{2(\phi+1)(\eta\phi+\eta+\phi)}P^2 \\ &- \frac{((2\phi+1)(\eta+\phi+\eta\phi-\tau(1+\phi))+\eta(\phi+1)^2\tau)}{(\phi+1)(\eta\phi+\eta+\phi)}P \\ &+ \frac{((\eta-\tau)(1+\phi)+\phi)^2}{2(\phi+1)(\eta\phi+\eta+\phi)} \end{aligned} \tag{3.15}$$

Conjecture the existence of a pure-strategy equilibrium in which case $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$. Substituting P^* into our expressions for V_L and V_1 implies:

$$\begin{aligned} H(\tau) = &\frac{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2}{2(1+\phi)(1+\eta+5\phi-\eta\phi)^2}\tau^2 \\ &- \frac{2\eta^2(1-\phi^2)+\phi(2+12\phi+7\phi^2)+\eta(1+9\phi+\phi^2-5\phi^3)}{(1+\phi)(1+\eta+5\phi-\eta\phi)^2}\tau \\ &+ \frac{\eta^3(\phi-1)^2(1+\phi)+\phi^2(2+11\phi)+2\eta\phi(1+8\phi+\phi^2)-4\eta^2\phi(-1+2\phi+2\phi^2)}{2(1+\phi)(1+\eta+5\phi-\eta\phi)^2} \end{aligned} \tag{3.16}$$

We have two possibilities to consider. First, suppose there is a pure-strategy equilibrium at the upper boundary of Case 1, $\tau = (2\eta - 1) \cdot P$. In this case, we know that $P = (\eta - \tau)/(\eta + 1)$ which implies $\tau = (2\eta - 1)/3$. We also know that $\phi \geq \frac{3\eta-2}{6-3\eta}$. Therefore, at the transition to this case we have $H\left(\frac{2\eta-1}{3}\right) = \frac{2-3\eta+\phi\cdot(6-3\eta)}{18(1+\phi)} \geq 0$ for $\phi \geq \frac{3\eta-2}{6-3\eta}$. Furthermore,

$$H'\left(\frac{2\eta-1}{3}\right) = -\frac{2+(10-6\eta)\phi+6(1-\eta)\phi^2}{3(1+\eta(1-\phi)+5\phi)(1+\phi)} < 0,$$

and

$$H''(\tau) = \frac{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2}{(1+\phi)(1+\eta+5\phi-\eta\phi)^2} > 0.$$

We see then that $H(\tau)$ is an upward facing parabola in τ. At the lower boundary of Case 2, where $\tau = \frac{2\eta-1}{3}$, $H(\tau)$ is positive but decreasing.

Solving $H(\tau) = 0$ yields two solutions:

$$\begin{aligned}
\tau_3, \tau_4 \quad &= \quad \frac{2\eta^2(1-\phi^2)+\eta(1+9\phi+\phi^2-5\phi^3)+2\phi+12\phi^2+7\phi^3}{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2} \\
&\pm \quad \frac{\sqrt{(1+\phi)(1+\eta+5\phi-\eta\phi)^2(\eta^2(1+\phi)(1+3\phi^2)-3\eta\phi^2(2+\phi)-2\phi^3)}}{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2}
\end{aligned}$$

If $\eta^2(1+\phi)(1+3\phi^2) - 3\eta\phi^2(2+\phi) - 2\phi^3 < 0$, the roots of the solution of $H(\tau) = 0$ are complex so $H(\tau) > 0$ for all τ. Therefore, for $\frac{2\eta-1}{3} < \tau \leq 1$, there is a pure-strategy equilibrium where all firms choose $F^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$ and $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$. However, our free disposal assumption implies $P^* \geq 0$ which requires $\tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$. We thus consider pure-strategy equilibrium in the range $\frac{2\eta-1}{3} < \tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$.

If $\eta^2(1+\phi)(1+3\phi^2) - 3\eta\phi^2(2+\phi) - 2\phi^3 \geq 0$, then τ_3 and τ_4 are real. Therefore, for $\tau_3 < \tau \leq \tau_4$, we have $H(\tau) < 0$ and a mixed-strategy equilibrium where P^* equates $V_L = V_1$. There are two solutions, but only one where P is less than the pure-strategy price (which must be true in equilibrium as argued above) and it is

$$\begin{aligned}
P^* \quad &= \quad \frac{\phi(1+2\phi(1-\tau)-3\tau)+\eta(1+\phi)(1+\tau+\phi(2+\tau))-\tau}{1+\phi(6-3\eta)(1+\phi)} \\
&- \quad \frac{\sqrt{(1+\phi)(\eta+\phi+\eta\phi)\begin{pmatrix} 3\eta^2\phi(1+\phi)-2(\phi(\tau-1)+\tau)^2 \\ +\eta(\phi(\tau-1)+\tau)(2+\phi(\tau-1)+\tau) \end{pmatrix}}}{1+\phi(6-3\eta)(1+\phi)}.
\end{aligned}$$

The fraction h of firms choosing $F = 1$ supporting the price P is found by equating demand for the risky asset $(1-h)\cdot(1-P-F_L)$ with supply $(1-h)\cdot(P+\phi F_L)/(1+\phi) + h\cdot(P+\phi\cdot 1)/(1+\phi)$ where $F_L = \frac{\tau(1+\phi)}{\eta+\phi+\eta\phi} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P$. Therefore, we have the unique proportion to clear the market:

$$h^* = \frac{(1+\phi)\cdot(\eta+\phi+\eta\phi-\tau(1+2\phi)-P(1+5\phi+\eta-\eta\phi))}{(1+2\phi)(\eta+\phi+\eta\phi-\tau-\phi\tau)-P(1+5\phi(1+\phi)-\eta\phi(1+\phi))}.$$

It has been verified numerically that $\tau_4 \leq 1$ and $P^*(\tau_4) > 0$. This means that for $\tau > \tau_4$, there will be a range of pure-strategy equilibria with positive prices. Therefore, for $\frac{2\eta-1}{3} < \tau \leq \tau_3$ and $\tau_4 < \tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$ we have a pure-strategy equilibrium where $F^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$ and $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$

Second, suppose there is a mixed-strategy equilibrium at the upper boundary of Case 1, i.e. $\tau_2 = (2\eta-1)\cdot P^*(\tau_2)$. Then we know that $\phi < \frac{3\eta-2}{6-3\eta}$. It can be verified that $H(\tau_2) < 0$, and, following the arguments above, $\{\tau_3, \tau_4\}$ are the same as described above. However, it can be verified that $\tau_3 < \tau_2$ if $\phi < \frac{3\eta-2}{6-3\eta}$, thus there is a mixed-strategy equilibrium for all $\tau_2 < \tau \leq \tau_4$. In the mixed-strategy equilibrium, $F_L = \frac{\tau(1+\phi)}{\eta+\phi+\eta\phi} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P^*$ and $\{P^*, h^*\}$ are characterized above. It is still the case that $\tau_4 \leq 1$ and $P^*(\tau_4) > 0$. Therefore, for $\tau_4 < \tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$ we have a pure-strategy equilibrium where $F^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}$ and $P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$

We must also check that all these equilibria fall under Case 2, which requires $\tau \leq \eta - P + (1-2P)\phi/(1+\phi)$. We know that $P^* \leq \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}$, the pure-strategy equilibrium price, for all equilibria. Therefore, it is sufficient to show that:

$$\tau \leq \eta - \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi} + \left(1 - \frac{2(\eta+\phi+\eta\phi-\tau(1+2\phi))}{1+\eta+5\phi-\eta\phi}\right)\frac{\phi}{1+\phi}$$

which is true if and only if

$$\tau \leq \frac{\eta+2\phi-\eta\phi}{1-\phi}.$$

However, this is satisfied for all $\tau \leq \frac{\eta+\phi+\eta\phi}{1+2\phi}$. Therefore, the equilibria described above all fall within Case 2.

Case 3: $\eta - P + (1-2P)\phi/(1+\phi) < \tau \leq 1$

If $\eta - P + (1-2P)\phi/(1+\phi) < \tau \leq 1$ then all firms choose $F > 1-P$ which implies that the aggregate demand for the liquidated asset is zero. But, since supply is positive when the price is positive, this case is incompatible with an equilibrium $P^* > 0$.

Equilibrium with $P^* = 0$

We assume free disposal therefore we know $P^* \geq 0$. We now consider the possibility that $P^* = 0$ which can occur in our model because we assume limited liability (i.e., the continuation value of a firm in distress is bounded below by zero). If $P = 0$ the four regions for the marginal cost of debt collapse into one:

- For $F \in [0, 1]$ the marginal cost of debt is $\eta \cdot F + F \cdot \phi/(1+\phi)$

Consequently, if $P = 0$ there can only exist a pure-strategy equilibrium in which all firms choose $F^* = \min\left\{1, \frac{\tau(1+\phi)}{\eta(1+\phi)+\phi}\right\}$. Therefore, the aggregate demand for the asset is $1 - P^* - F^* = 1 - F^* \geq 0$. The aggregate supply of the asset, however, is indeterminate. In particular, because of limited liability the firm will be indifferent between liquidation at $P^* = 0$ and continuation for all $V_i \in \left[0, \frac{\phi F^*}{1+\phi}\right]$. Therefore, the price $P^* = 0$ can be supported in equilibrium if $1 - F^* \leq \frac{\phi F^*}{1+\phi}$ or $\tau \geq \frac{\eta(1+\phi)+\phi}{1+2\phi}$.

Proof when $\eta < 1/2$

We now consider the case $\eta < 1/2$. For now, assume $P > 0$. We will consider the possibility that $P = 0$ at the end of the proof. If $P > 0$ the marginal cost of debt now falls into four regions:

1. For $F \in [0, P]$ the marginal cost of debt is 0
2. For $F \in [P, P \cdot (1/\eta - 1)]$ the marginal cost of debt is $(F - P)\phi/(1+\phi)$
3. For $F \in [P \cdot (1/\eta - 1), 1 - P]$ the marginal cost of debt is $\eta F - P(1-\eta) + (F-P)\phi/(1+\phi)$

4. For $F \in [1-P, 1]$ the marginal cost of debt is $(F-P)\phi/(1+\phi)$

We must consider three cases: (i) $0 \leq \tau < (P/\eta - 2P)\phi/(1+\phi)$, (ii) $(P/\eta-2P)\phi/(1+\phi) \leq \tau \leq \eta - P + (1-2P)\phi/(1+\phi)$, and (iii) $\eta - P + (1-2P)\phi/(1+\phi < \tau \leq 1$.

Case 1: $0 \leq \tau < (P/\eta - 2P)\phi/(1+\phi)$

If $0 \leq \tau < (P/\eta - 2P)\phi/(1+\phi)$ then $F \in [P, P(1/\eta - 1)]$ and all firms equate the marginal cost of debt in this region to the marginal benefit which implies

$$F(P) = P + \frac{\tau(1+\phi)}{\phi}.$$

In this region, all firms for whom the asset is positive NPV ($\eta V \geq P$) will be able to obtain financing to purchase the asset. The demand for the liquidated asset is then $1 - P/\eta$ and the supply of the liquidated asset is $(P + \phi F)/(1+\phi) = P + \tau$. Equating supply and demand gives the equilibrium price

$$P^* = \frac{\eta(1-\tau)}{1+\eta}$$

and, therefore,

$$F^* = \frac{\tau(1+\eta+\phi) + \phi\eta}{\phi(1+\eta)}.$$

To determine the values of τ included in this case, substitute P^* into the expression

$$\tau < (P/\eta - 2P)\phi/(1+\phi) \quad \Rightarrow \quad \tau \leq \frac{\phi(1-2\eta)}{(1+\eta)(1+\phi)+\phi(1-2\eta)} \equiv \tau_0.$$

Case 2: $(P/\eta - 2P)\phi/(1+\phi) \leq \tau \leq \eta - P + (1-2P)\phi/(1+\phi)$

If $(P/\eta - 2P)\phi/(1+\phi) \leq \tau \leq \eta - P + (1-2P)\phi/(1+\phi)$ then firms choose either $F_L \in [P(1/\eta - 1), 1-P)$ or $F_H \in [1-P, 1]$ where $F_L = \frac{\tau(1+\phi)}{\eta+\phi+\eta\phi} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)}P$ and $F_H = \min\{1, P + \tau\cdot(1+\phi)/\phi\}$.

<u>Pure-strategy equilibria</u>

There cannot exist a pure-strategy equilibrium with $P > 0$ in which all firms choose F_H because the aggregate demand for the risky asset would be zero but the supply is positive

$[(P+\phi F)/(1+\phi)]$. Therefore, in a pure-strategy equilibrium, firms choose

$$F = \frac{\tau(1+\phi)}{\eta(1+\phi)+\phi} + \frac{(1-\eta)(1+\phi)+\phi}{\eta(1+\phi)+\phi} P.$$

The demand for the liquidated asset is $1-P-F$ and the supply of the liquidated asset is

$$\Lambda = \frac{P+\phi F}{1+\phi} = \frac{P \cdot [(1-\phi)\eta + 2\phi] + \phi\tau}{\eta(1+\phi)+\phi}.$$

Equating supply and demand gives the equilibrium price

$$P^* = \frac{\eta+\phi+\eta\phi-\tau(1+2\phi)}{1+\eta+5\phi-\eta\phi}.$$

Substituting into the expression for F yields

$$F^* = \frac{1+2\phi+\tau+(\tau-\eta)(1+\phi)}{1+\eta+5\phi-\eta\phi}.$$

Mixed-strategy equilibria

Again, there is the possibility of a mixed-strategy equilibrium. The proof here follows closely the proof in Case 2 when $\eta \geq 1/2$. Firms choosing $F_L = \frac{\tau(1+\phi)}{\eta+\phi+\eta\phi} + \frac{(1-\eta)(1+\phi)+\phi}{(\eta+\phi+\eta\phi)} P$ have ex ante value V_L as described in equation (3.12), firms choosing $F_H = P + \tau \cdot (1+\phi)/\phi$ have ex ante value V_H as described in equation (3.13), and firms choosing $F_H = 1$ have ex ante value V_1 as described in equation (3.14).

It is straightforward to show that Lemma 2 applies in the case $\eta < 1/2$ when $\tau_0 \leq \tau < \phi/(1+\phi)$. Therefore, in a mixed-strategy equilibrium the high type always chooses $F_H = 1$. Therefore, we must only compare V_L to V_1 to find a mixed-strategy equilibrium. We also have $V_L - V_1$ as described in equation (3.15) and $H(\tau)$ as described in equation (3.16).

We know there is a pure-strategy equilibrium at the upper boundary of case 1, $\tau = \tau_0$. Therefore, at the transition to this region we have $H(\tau_0) \geq 0$. Furthermore, it can be shown numerically that for all $\phi \in [0,1]$ and all $\eta \in [0, 1/2]$ that

$$H'(\tau_0) < 0,$$

and

$$H''(\tau) = \frac{2+14\phi+20\phi^2+9\phi^3+3\eta(1-\phi)(1+\phi)^2}{(1+\phi)(1+\eta+5\phi-\eta\phi)^2} > 0.$$

We see then that $H(\tau)$ is an upward facing parabola in τ. At the lower boundary of Case 2, where $\tau = \tau_0$, $H(\tau_0)$ is positive but decreasing.

Solving $H(\tau) = 0$ yields τ_3, τ_4 as before and the remainder of the proof is identical to Case 2 when $\eta \geq 1/2$.

Case 3: $\eta - P + (1 - 2P)\phi/(1 + \phi) < \tau \leq 1$

If $\eta - P + (1 - 2P)\phi/(1 + \phi) < \tau \leq 1$ then all firms choose $F > 1 - P$ which implies that the aggregate demand for the liquidated asset is zero. But, since supply is positive when the price is positive, this case is incompatible with an equilibrium $P^* > 0$.

Equilibrium with $P^* = 0$

Finally, as before, if $P = 0$ the four regions for the marginal cost of debt collapse into one:

- For $F \in [0, 1]$ the marginal cost of debt is $\eta \cdot F + F \cdot \phi/(1 + \phi)$

Consequently, if $P = 0$ there can only exist a pure-strategy equilibrium in which all firms choose $F^* = \min\left\{1, \frac{\tau(1+\phi)}{\eta(1+\phi)+\phi}\right\}$. Following the argument in the proof when $\eta \geq 1/2$, the price $P^* = 0$ can be supported in equilibrium if $\tau \geq \frac{\eta(1+\phi)+\phi}{1+2\phi}$.

Uniqueness

We've established above that the equilibrium debt choices, prices, and high-type proportion are unique functions of the exogenous parameters in each of the equilibrium types (pure strategy interior/extreme and mixed-strategy). Also, the necessary conditions for each of the equlibria form non-overlapping regions. It cannot be that a given exogenous parameter value supports multiple types of symmetric equilibria. Therefore, the equilibium is unique for any given parameter values.

3.8 Appendix: Extension to Immediate Use of Debt Benefits

We now show that the solution to our model is isomorphic to one in which the firm can use the immediate benefit τF to pay off debt and purchase liquidated assets.

If $F(1-\tau) < P$ the value of the firm is now

$$\int_0^P P dV + \int_P^1 V dV + \int_{P+F(1-\tau)}^1 \max\{0, \eta V - P\} dV + \tau F.$$

The supply of the risky asset is P and the demand is now $1 - P - (1-\tau)F$.

Therefore, if we let $F' = (1-\tau)F$ the value of the firm is

$$\int_0^P P dV + \int_P^1 V dV + \int_{P+F'}^1 \max\{0, \eta V - P\} dV + \frac{\tau}{1-\tau} \cdot F'.$$

The supply of the risky asset is P and the demand is now $1 - P - F'$.

If $F(1-\tau) \geq P$ the value of the firm is

$$\int_0^{\Lambda'} P dV + \int_{\Lambda'}^{F(1-\tau)} [V - \phi(F - (V + \tau F)] dV + \int_{(1-\tau)F}^1 V dV + \int_{P+F(1-\tau)}^1 \max\{0, \eta V - P\} dV + \tau F.$$

The supply of the risky asset is $\Lambda' = \frac{P+\phi(1-\tau)F}{1+\phi}$ and the demand is now $1 - P - (1-\tau)F$.

Again, if we let $F' = (1-\tau)F$ the value of the firm is

$$\int_0^{\Lambda'} P dV + \int_{\Lambda'}^{F'} [V - \phi(F' - V)] dV + \int_{F'}^1 V dV + \int_{P+F'}^1 \max\{0, \eta V - P\} dV + \frac{\tau}{1-\tau} \cdot F'.$$

The supply of the risky asset is $\Lambda' = \frac{P+\phi F'}{1+\phi}$ and the demand is now $1 - P - F'$.

In sum, the solution to our original model yields $\{P(\tau'), F'(\tau')\}$ where $\tau' = \tau/(1-\tau)$. To convert to the equilibrium $\{P(\tau), F(\tau)\}$ note that $\tau = \tau'/(1+\tau')$ and $F = F'/(1-\tau)$. The latter expression for the face value of debt implies that the comparative statics for debt choices with respect to model parameters are not necessarily the same as in the base model. Although we don't report the results here, our main qualitative results continue to hold in this extension; in particular, debt levels and ratios may increase or decrease in η and the comparative statics for debt levels can be different than for debt-to-value ratios.

3.9 Appendix: Extension to Industry Booms and Busts

In this section we extend our model to consider uncertainty in the economy. For certain parameter regions of this model, it is possible that we find either pro-cyclical or counter-cyclical reallocations of assets.

Suppose we extend the model to include uncertainty about the distribution of asset qualities, $v_i \in [0, \gamma]$. With probability a, $\gamma = 1-\Delta$ and with probability $1-a$, $\gamma = 1+\Delta$. In this model, the firms will choose their debt level F and mixture probability h before realizing the state of the economy, but $P(\gamma)$ will be determined after the state of the economy has been realized.

3.9.1 Reallocation

We know that the amount of reallocation will equal the mass of firms who are acquiring assets. This is equal to:

$$Q(\gamma) = (1-h) \cdot \min\left\{\frac{\gamma - P(\gamma) - F}{\gamma}, \frac{\gamma - P(\gamma)/\eta}{\gamma}\right\}$$

In our base model, the marginal cost of debt may fall in several regions depending on the debt level. However, the boundaries of these regions depend on the realization of γ. Thus, optimal debt levels will be determined by equating the marginal benefit of debt, τ, with the expected marginal cost. We will assume for simplicity that Δ is sufficiently small such that the optimal debt level falls in the same marginal cost region in either case (i.e. there are not distress costs in the bust or the boom.)

3.9.2 Pro-Cyclical

Paralleling Case 1 where $\eta \geq 1/2$ of Internet Appendix A, in this region we have no distress costs, and our only concern is foregone acquisition costs. Conditional on γ, we know that the supply is $\frac{P(\gamma)}{\gamma}$ and the demand is $\frac{\gamma - P(\gamma) - F}{\gamma}$. Therefore, $P(\gamma) = \frac{\gamma - F}{2}$. Plugging this back in to our supply we find that reallocation here is $Q(\gamma) = \frac{1}{2} - \frac{F}{2\gamma}$ and it is the case that $Q(1+\Delta) > Q(1-\Delta)$. We have in this region, the case where booms lead to more reallocation.

For this equilibrium to exist, there must exist parameters such that $F < P(\gamma)$ for both values of γ, which has been verified to exist numerically.

3.9.3 Counter-Cyclical

Paralleling Case 1 where $\eta < 1/2$ of Internet Appendix A, in this region we have only distress costs and no foregone acquisition costs (still acquisitions, just NPV constraint binds). Conditional on γ, we know that the supply is $\frac{P(\gamma)+\phi F}{\gamma(1+\phi)}$ and the demand is $\frac{\gamma - P(\gamma)/\eta}{\gamma}$. Therefore, $P(\gamma) = \frac{\eta(\gamma(1+\phi)-\phi F)}{1+\eta+\phi}$. Plugging this back in to our demand we find that reallocation here is $Q(\gamma) = 1 - \frac{1+\phi}{1+\eta+\phi} + \frac{\phi F}{\gamma(1+\eta+\phi)}$ and it is the case that $Q(1+\Delta) < Q(1-\Delta)$. We have in this region the case where booms lead to less reallocation.

Again, for this equilibrium to exist, there must exist parameters such that $F < P(\gamma)(1/\eta - 1)$ for both values of γ, which has been verified to exist numerically.

3.10 Appendix: Tax Shields as Sources of Debt Benefits

In this section, we allow the debt benefits to be partly or fully tax-related. t is now the total benefit of debt, inclusive of the direct benefits (e.g. signaling or agency related) and the tax shield benefits. $g \in [0,1]$ is the share of the debt benefits that are from tax shields. The value of the firm for a given level of debt is

$$V(P, F_i) = t \cdot F_i + (1 - g \cdot t) \tag{3.17}$$

$$\cdot \begin{cases} \int_0^P P\,dv + \int_P^1 v\,dv + \int_{\min\{P+F_i,1\}}^1 \max\{0, \eta \cdot v - P\}\,dv & \text{if } F_i < P \\ \int_0^{\Lambda(F_i)} P\,dv + \int_{\Lambda(F_i)}^{F_i} [v - \phi \cdot (F_i - v)]\,dv + \int_{F_i}^1 v\,dv + \int_{\min\{P+F_i,1\}}^1 \max\{0, \eta \cdot v - P\}\,dv & \text{if } F_i \geq P \end{cases} \tag{3.18}$$

The tax rate on the total earnings of the firm is $g \cdot t$, and the total benefits of debt are $t \cdot F_i$.

This means that the tax shield is $g \cdot t \cdot F_i$, and other direct debt benefits are $(1-g) \cdot t \cdot F_i$. If $g = 0$, it is the model in the main text. If $g = 1$, all of the benefits of debt come from the tax shield. Our specification allows full deductibility of the debt expense, in all states–possibly by selling the tax-loss credits. It also assumes that the proceeds in liquidation, P, are fully taxable. This is equivalent to the asset being fully depreciated with a tax basis of zero.

Dividing equation (3.17) by $(1 - g \cdot t)$ returns to our original model with $\tau = t/(1 - g \cdot t)$. However, the "firm value" in our original model is now $\tilde{V} = V/(1 - g \cdot t)$. The optimal F_i^* is the same for either problem (with the remapped τ), because dividing by a constant does not change the optimum. Also, the rescaling of firm value does not affect any other equilibrium quantities such as P^* or h^*. However, we only solved the original model for $\tau \leq 1$. To preserve this, we also require that $t/(1 - g \cdot t) \leq 1$ or $t \leq 1/(1 + g)$.

The sign of most of our comparative statics are preserved, but the quantitative plots (numbers) have to be remapped. For comparative statics with respect to η and ϕ, the comparative statics are the same as the old model with the remapping $\tau = t/(1 - g \cdot t)$. The comparative statics with respect to the new argument t can be found with the chain rule where $\tau(t) = t/(1 - g \cdot t)$:

$$\frac{\partial F(\tau(t))}{\partial t} = \frac{\partial F(\tau(t))}{\partial \tau} \cdot \frac{\partial \tau(t)}{\partial t} = \frac{\partial F(\tau(t))}{\partial \tau} \cdot \frac{1}{(1 - g \cdot t)^2}$$

We have $\frac{\partial F(\tau(t))}{\partial \tau}$ from our original model, so the sign is preserved. It is still the case that the debt face value is increasing in t. The same chain rule argument works for our other equilibrium quantities, *except* where we use V^* explicitly. This is V^* and D^*/V^*.

The $\partial V^*/\partial t$, where $V^* \equiv \tilde{V}^* \cdot (1 - g \cdot t)$, is

$$\begin{aligned} \frac{\partial V^*}{\partial t} &= (1 - g \cdot t) \cdot \frac{\partial \tau(t))}{\partial t} \cdot \frac{\partial \tilde{V}^*(\tau(t))}{\partial \tau} \; - \; g \cdot \tilde{V}^*(\tau(t)) \\ &= \frac{1}{1 - g \cdot t} \cdot \frac{\partial \tilde{V}^*(\tau(t))}{\partial \tau} \; - \; g \cdot \tilde{V}^*(\tau(t)) \end{aligned}$$

Although $\partial \tilde{V}^*/\partial \tau$ in our original model is always positive, $\partial V^*/\partial t$ is *not* generally positive, especially when $g \approx 1$. Instead, and somewhat surprisingly, $\partial V^*/\partial t$ becomes ambiguous: It is now negative for low tax rates, but $\partial V^*/\partial t$ is still positive for high tax rates in equilibrium. The top left of Figure 3.10 shows that this is not an obscure region, but a widespread phenomenon for high tax rates. This is mostly due to the fact that, in this region, the face value of debt exceeds the expected EBIT of the firm. The firm expects to receive more in tax-loss credits than it expects to have to pay out in taxes.

With even V^* being ambiguous in t, it is a lesser surprise that D^*/V^* also retains the ambiguity in both t and τ. This can be seen as follows: Let what we have be $a(\tau) = D^*(\tau)/\tilde{V}^*(\tau)$ and what we want be $b(t) = D^*(\tau(t))/V^*(t) = D^*(\tau(t))/(\tilde{V}^*(\tau(t))(1-g\cdot t)) = a(\tau(t)) \cdot 1/(1-g\cdot t)$. Then

$$\begin{aligned}\frac{\partial D^*(\tau(t))/V^*(\tau(t))}{\partial t} = \frac{\partial b}{\partial t} &= a'(\tau(t)) \cdot \frac{\partial \tau(t)}{\partial t} \cdot \frac{1}{1-g\cdot t} + a(\tau(t)) \cdot \frac{g}{(1-g\cdot t)^2} \\ &= a'(\tau(t)) \cdot \frac{1}{(1-g\cdot t)^3} + a(\tau(t)) \cdot \frac{g}{(1-g\cdot t)^2}\end{aligned}$$

Thus, D^*/V^* is not necessarily decreasing in actual tax rates, either.

In sum, our model retains all the same comparative statics, regardless of the source of the debt benefits, except where we discuss it in our paper—specifically, in the aforementioned $\partial V^*/\partial\tau$ case. (Although the quantitative regions can also change, none change so dramatic as to undo or now deserve a "rare" designation.)

Figure 3.1: Game Tree for the Acquisition Model

Figure 3.1 illustrates the timing of managers' choices and the resulting payoffs in the acquisition model. The manager first chooses a debt level. The value of the asset is then realized, and the manager decides whether or not to sell the asset and whether or not to acquire another. The assets then pay off according to the previous decisions.

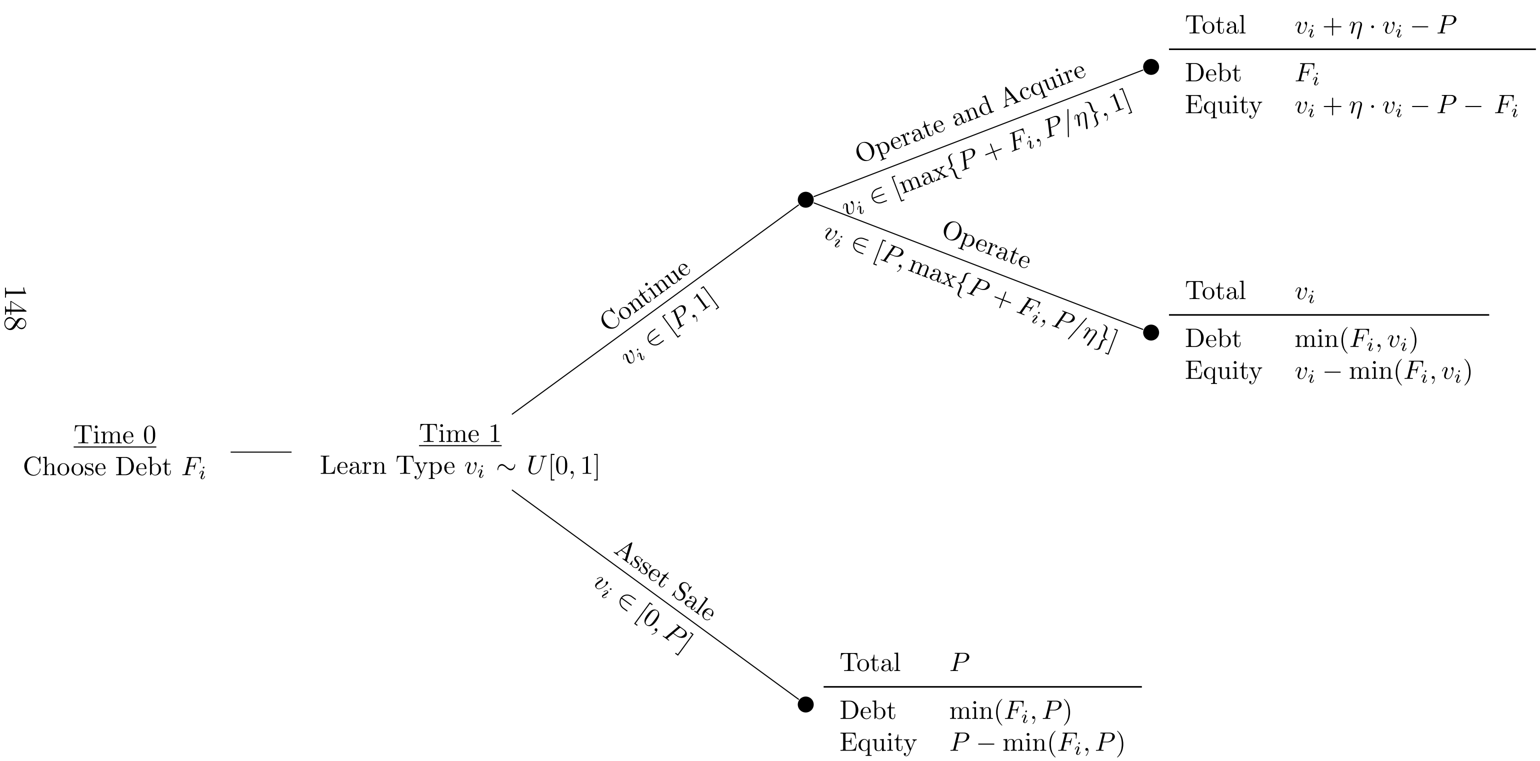

Figure 3.2: Comparative Statics for Heterogeneity in the Acquisition-Only Model ($\phi = 0$)
Figure 3.2 shows a contourplot with the fraction of heterogeneous firms as the dependent variable. The yellow area contains the two-type equilibria (as defined in Theorem 7 on Page 96). The area above the diagonal is uninteresting, as all firms choose $F_i = 1$ and the price is 0. The area on the bottom right has all firms act alike. Heterogeneity arises unless redeployability is high and debt benefits are low. It is common for intermediate values of redeployabilities and direct debt benefits.

Frequency of Max-Debt Types, h^*

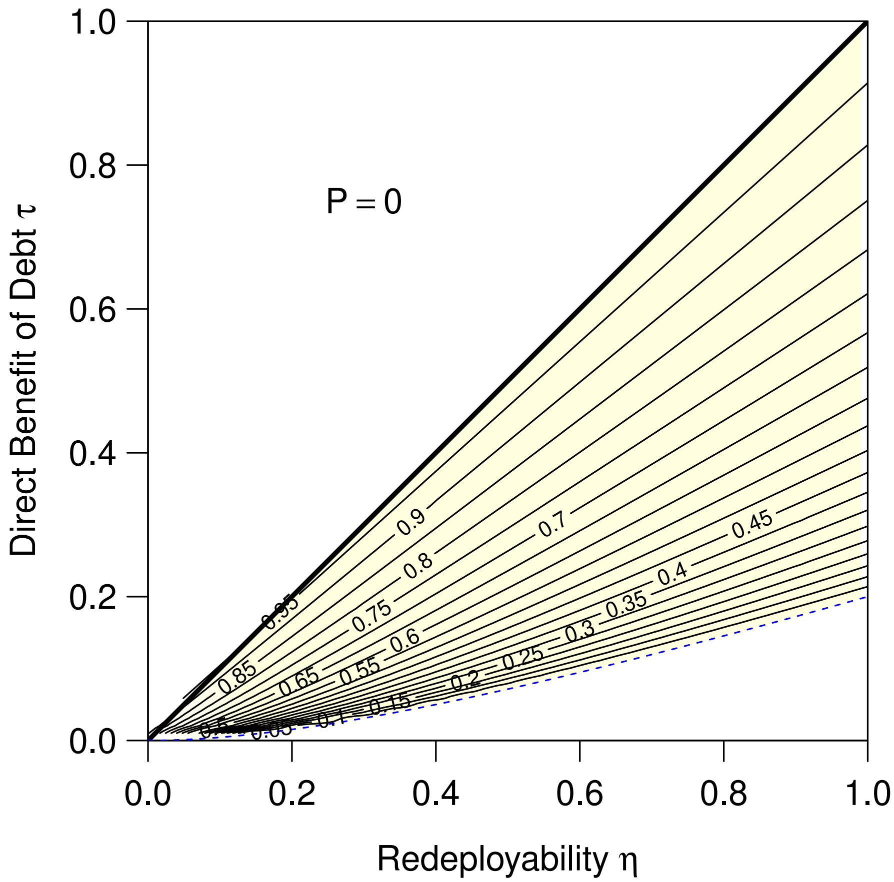

Figure 3.3: Comparative Statics for the Acquisition-Only Model Leverage (ϕ=0)

Figure 3.3 shows contourplots for the acquisition-only model. The yellow area contains the two-type equilibria (as defined in Theorem 7 on Page 96). Patterns that have "∩" or "∪" shapes indicate ambiguous comparative statics in redeployability. "⊂" or "⊃" shapes indicate ambiguous comparative statics in direct debt benefits. Left and Middle: The face value F^* and current market value of debt (D^*) increase [everywhere for the industry, almost everywhere for the low-debt firm] monotonically in debt benefits τ and decrease monotonically in redeployability η. Right: The debt-to-value ratio is ambiguous in debt benefits τ, and decreasing monotonically in redeployability η.

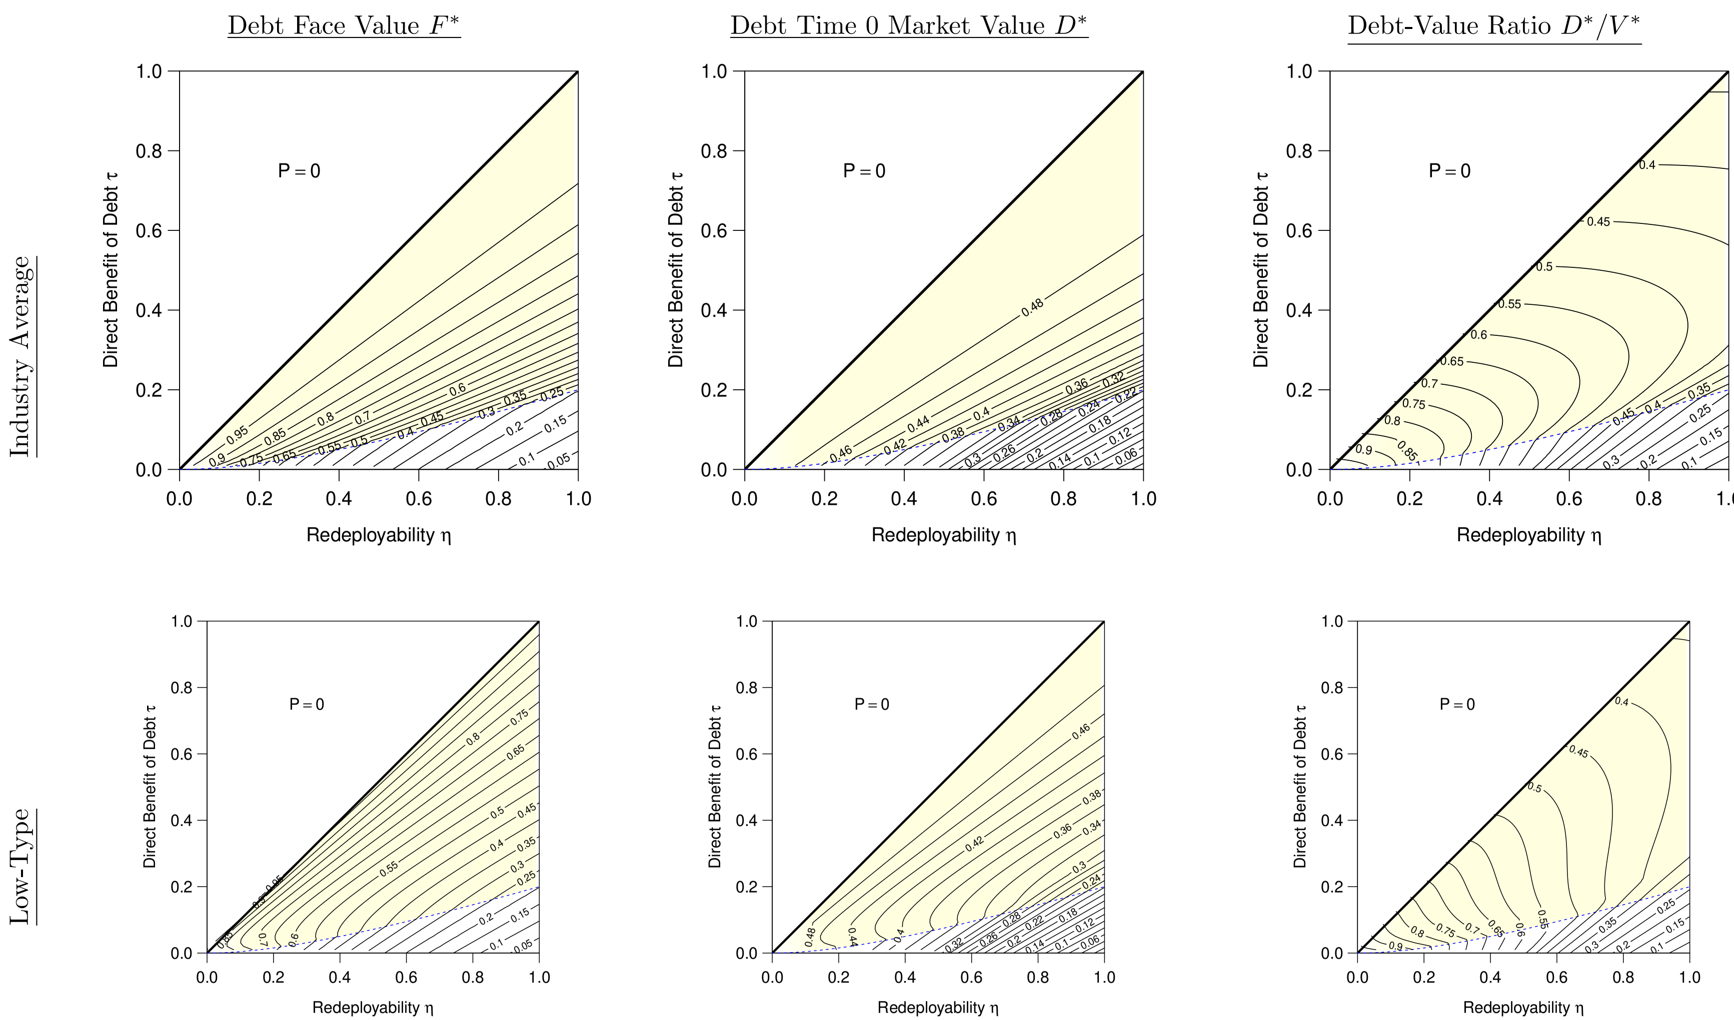

Figure 3.4: Peer Effects on Debt Choice

Figure 3.4 shows equilibrium prices and debt choices when η=1/2. For high τ, some firms choose a high-debt strategy (F_H=1). Therefore, industry debt (F_{Ind}) is higher and the equilibrium price (P) is lower than what would have occurred if all firms had chosen a low-debt strategy (represented by the dashed lines). Other firms recognize that more valuable buying opportunities will become available and have an incentive to choose debt (F_L) below what is optimal if industry debt was lower.

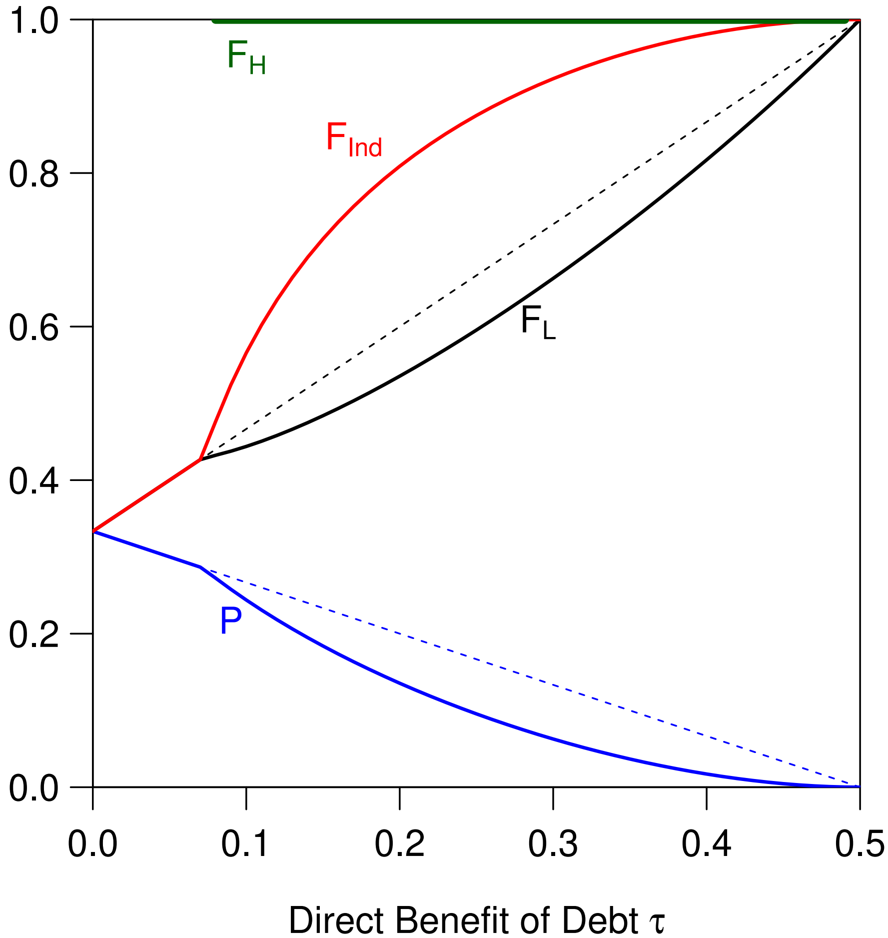

Figure 3.5: Game Tree for the Full Model ($F_i > P$)

Figure 3.5 illustrates the timing of manager's choices and the resulting payoffs in the full model. The manager first chooses a debt level. The value of the asset is then realized, and the manager decides whether or not to default and whether or not to acquire another asset. In the event of default, the manager also chooses between liquidation and reoganization. The assets then pay off according to the previous decisions.

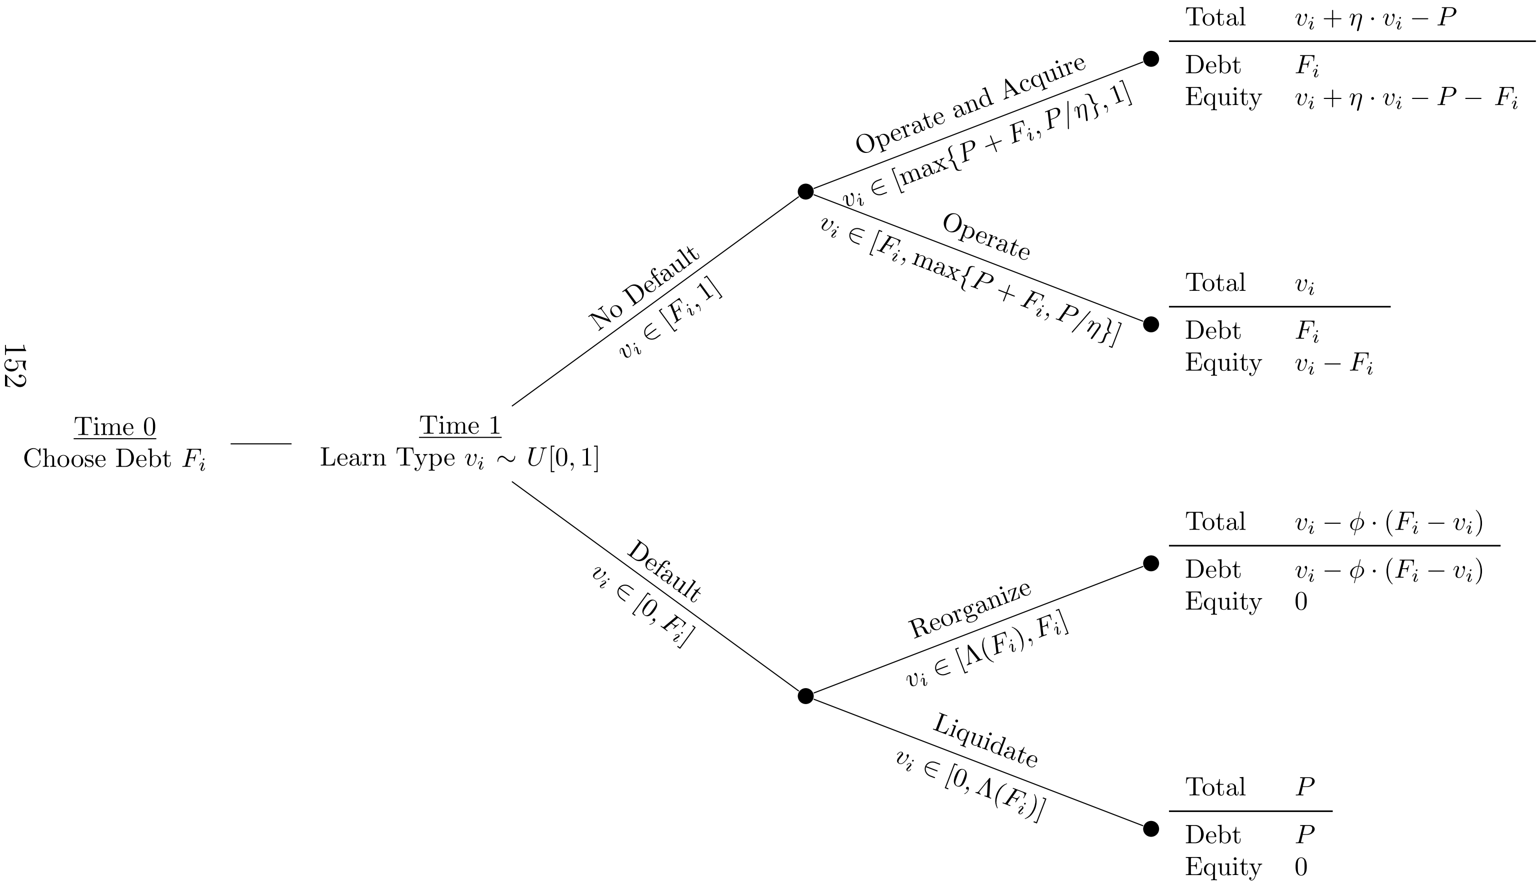

Figure 3.6: Comparative Statics for Heterogeneity when ϕ=0.25

Figure 3.6 shows a contourplot with the fraction of heterogeneous firms as the dependent variable in the full model with distress costs. The yellow area contains the two-type equilibria (as defined in Theorem 8 and Internet Appendix A). Patterns that have "∩" or "∪" shapes indicate ambiguous comparative statics in redeployability. "⊂" or "⊃" shapes indicate ambiguous comparative statics in direct debt benefits. Heterogeneity still arises for intermediate values of debt benefits and large values of redeployability. However, the heterogeneity region is now smaller than it was when ϕ=0 in Figure 3.2.

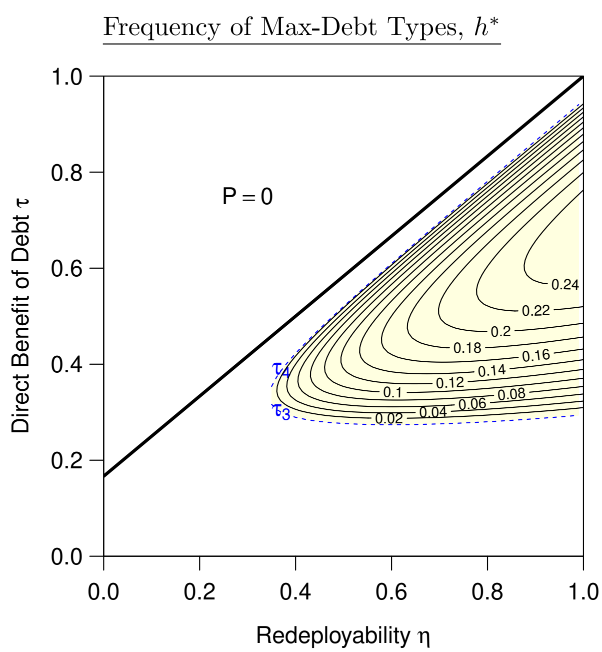

Figure 3.7: Comparative Statics for Industry Leverage when ϕ=0.25

Figure 3.7 shows contourplots for the full model with distress costs. The yellow area contains the two-type equilibria (as defined in Theorem 8 and Internet Appendix A). Patterns that have "∩" or "∪" shapes indicate ambiguous comparative statics in redeployability. "⊂" or "⊃" shapes indicate ambiguous comparative statics in direct debt benefits. Note somewhere $F^*_{\text{Ind}} = h^* \cdot 1 + (1 - h^*) \cdot F^*_L$. — — Industry debt increases monotonically with direct debt benefits τ. However, it can first increases and then decrease in redeployability η. Not shown, the low-debt firm (rather than industry debt) shows a very similar pattern, with only small changes in the yellow region.

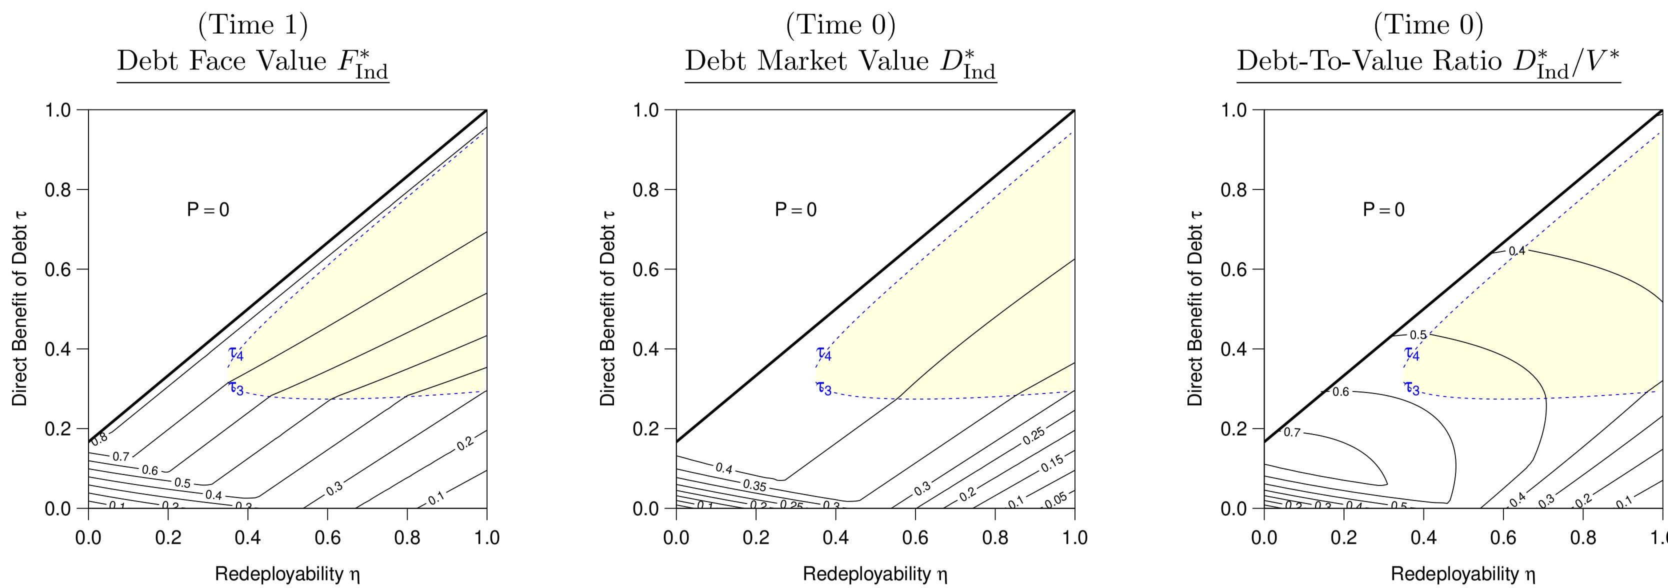

Figure 3.8: Ancillary Comparative Statics for ϕ=0.25

Figure 3.8 shows contourplots for the full model with distress costs. The yellow area contains the two-type equilibria (as defined in Theorem 8 and Internet Appendix A). Patterns that have "$\cap$" or "$\cup$" shapes indicate ambiguous comparative statics in redeployability. "$\subset$" or "$\supset$" shapes indicate ambiguous comparative statics in direct debt benefits.

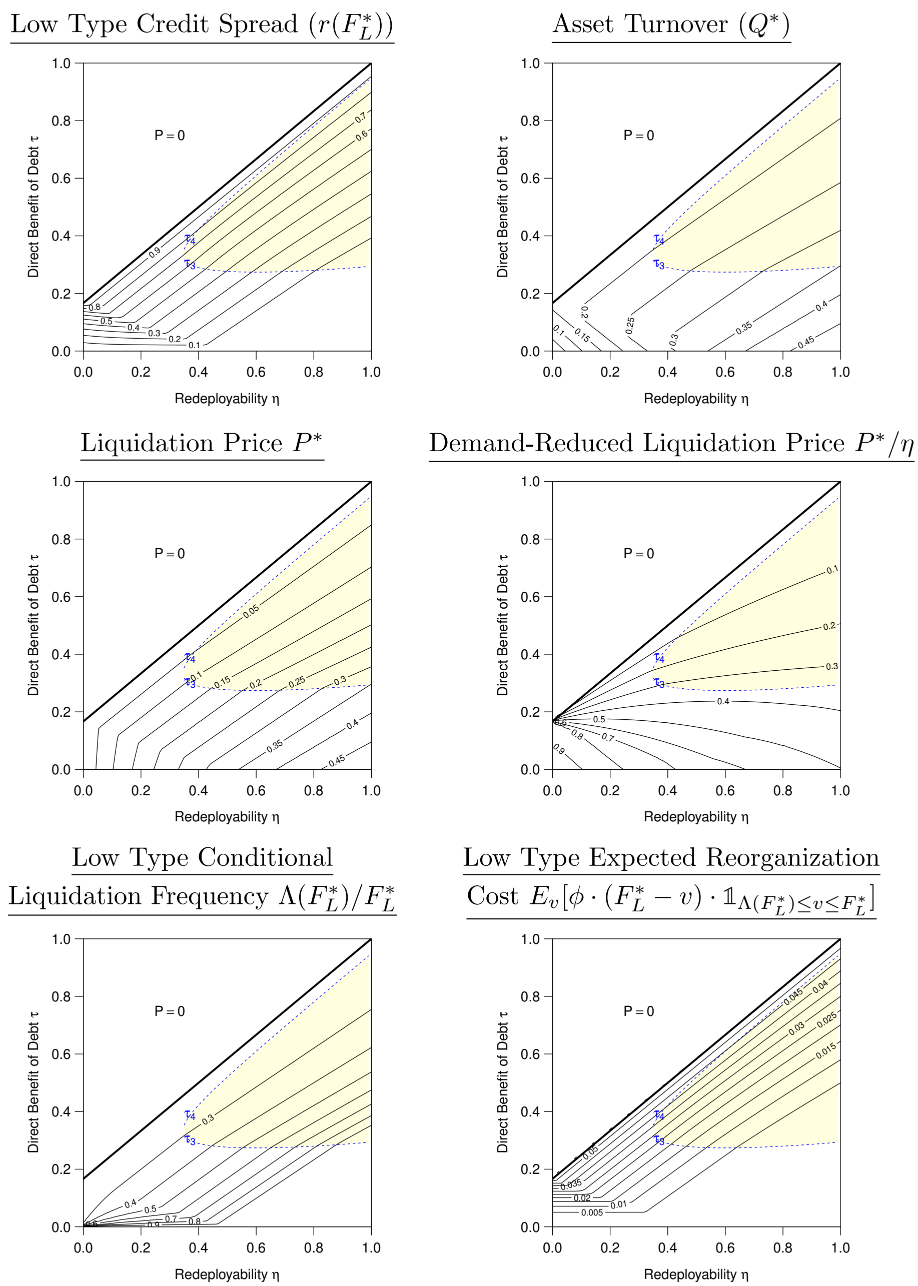

Figure 3.9: Allocational Efficiency

Figure 3.9 shows contourplots for varying levels of distress costs. The yellow area contains the two-type equilibria (as defined in Theorem 8 and Internet Appendix A). The fat line shows the parameters for τ and η where equilibrium results in first-best redeployment. The area to the left of the fat line has too much transfer activity (Q^*). The area to the right of the fat line has too little transfer activity. If there are no reorganization costs (ϕ=0), there is always too little redeployment.

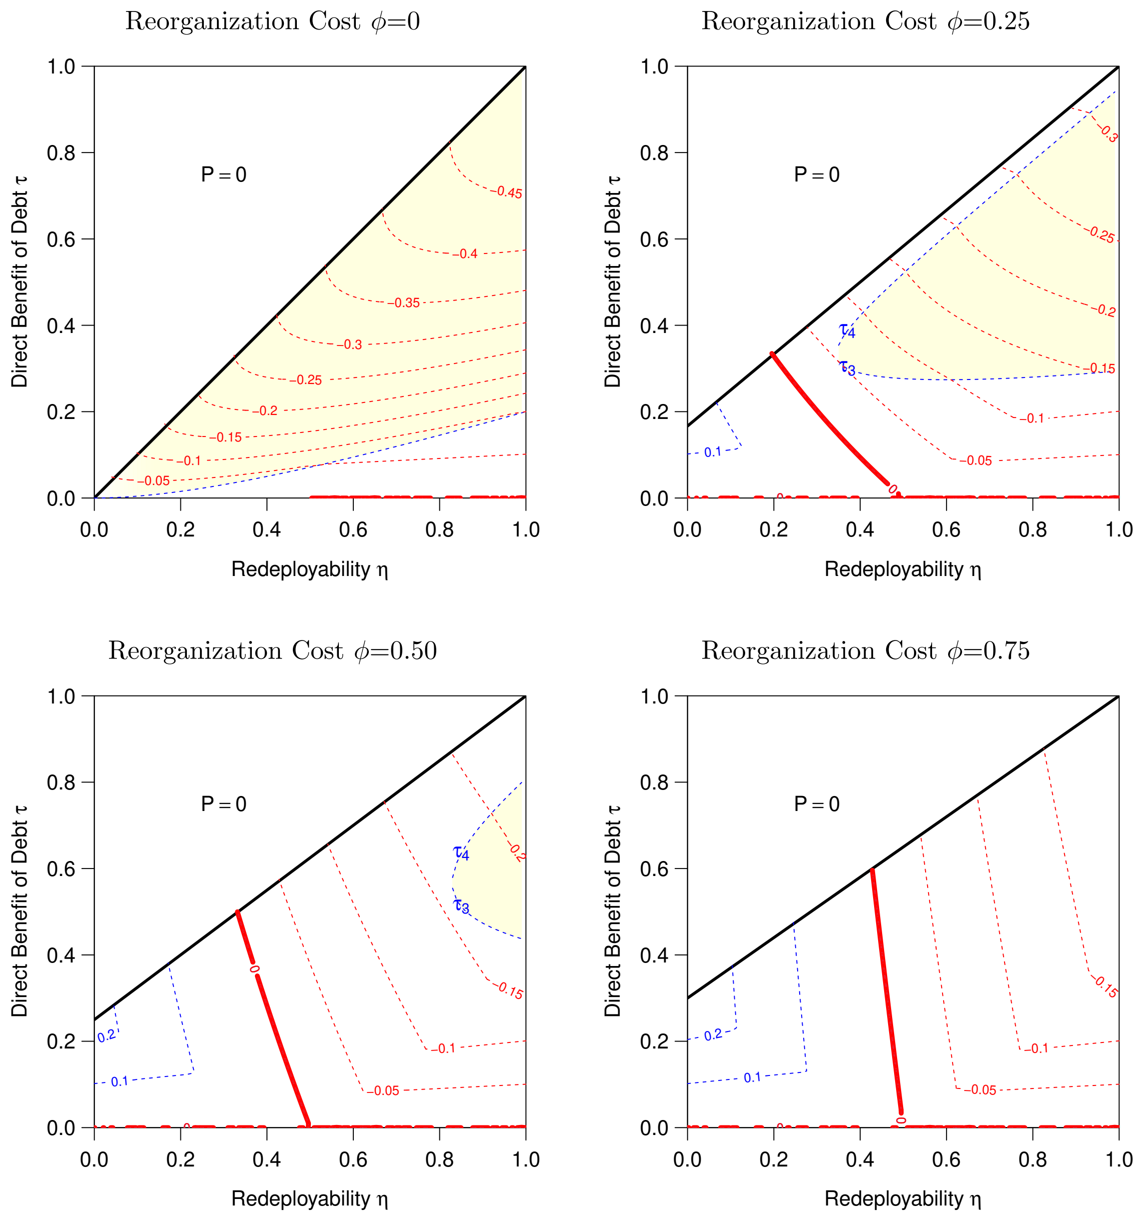

Figure 3.10: Value and Debt-Ratio When Benefits are Tax Shields

Figure 3.10 shows contourplots where the benefits of debt come from tax shields vs. when they come from other direct benefits. When there are taxes and debt benefits derive from the tax shield, value becomes ambiguous in the tax rate.

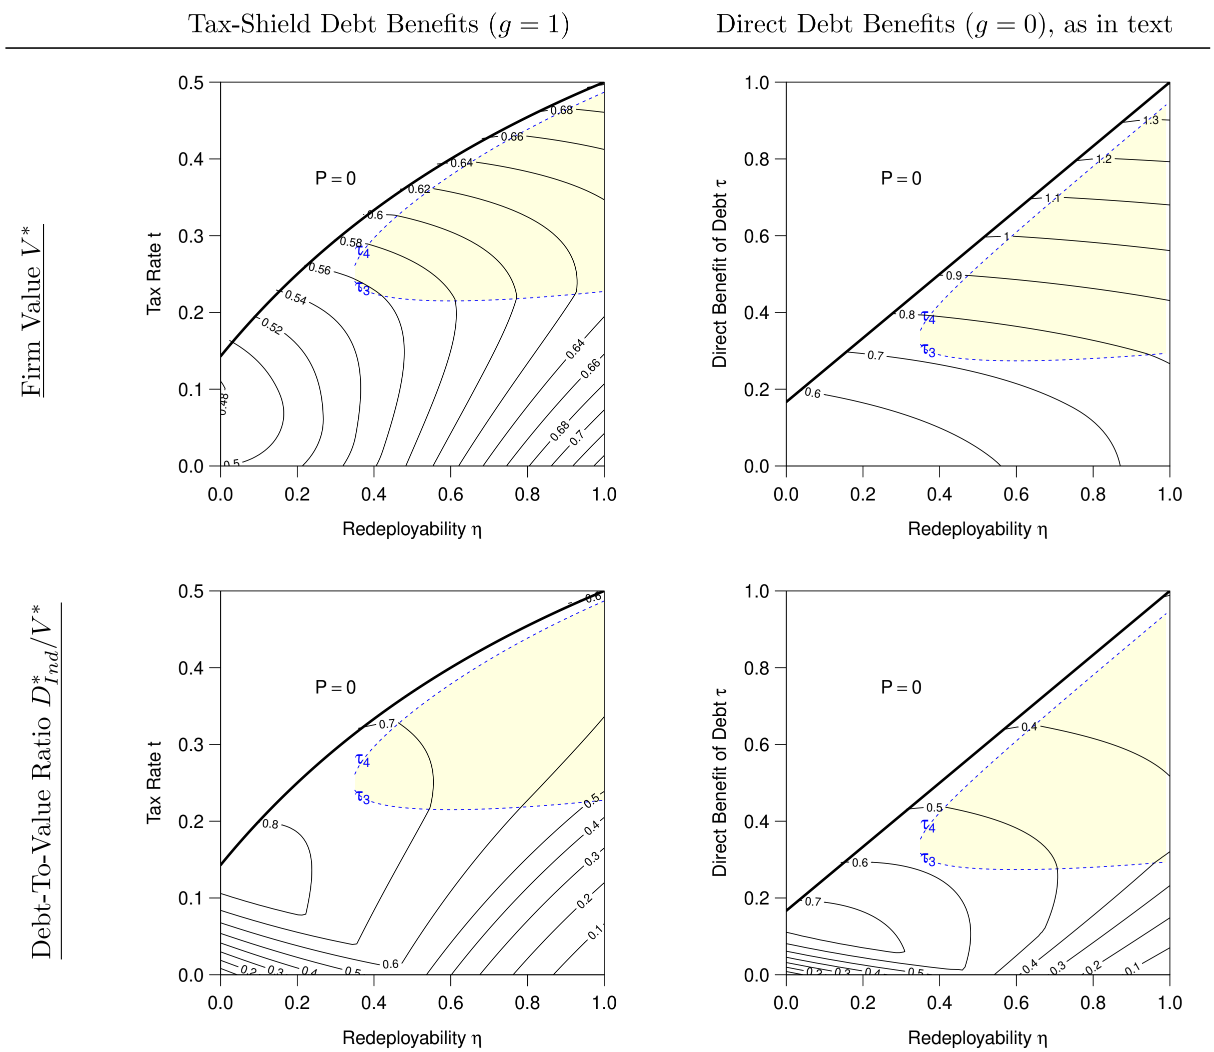

Table 3.1: Variables

Table 3.1 defines the variables used throughout the model.

v_i	$v_i \sim U[0,1]$	Unlevered firm/asset type
<u>Exogenous Parameters</u>		
ϕ	$0 \leq \phi \leq 1$	Reorganization impairment $\phi \cdot (F_i - v_i)$ for firms continuing in default.
η	$0 \leq \eta \leq 1$	Asset redeployability
τ	$0 \leq \tau \leq 1$	Other (net) benefits of debt
<u>Endogenous Quantities</u>		
F_i	$0 \leq F_i \leq 1$	Face Value of Debt for firm i, promised **for time 1**.
D_i	$0 \leq D_i \leq F_i$	Value of Debt **at time 0**, as in (3.9)
h	$0 \leq h \leq 1$	Proportion of F_H^*=1 types
$\Lambda(F_i, P)$	$0 \leq \Lambda \leq 1$	Liquidation/continuation threshold, $(P + \phi \cdot F_i)/(1+\phi)$
$V(F_i, P)$	$V \geq 0$	Firm value **at time 0**
P	$P \geq 0$	Price of liquidated assets **at time 1**
Q	$Q \geq 0$	Assets transferred **at time 1**
r	$r \geq 0$	Credit Spread, $F_i^*/D_i^* - 1$, as in (3.9)

Table 3.2: Summary of Comparative Statics

Table 3.2 describes the comparative statics of endogenous variables with respect to the exogenous parameters for the full model with distress costs. Ambiguous comparative statics are illustrated with two examples (order η, ϕ, τ), in which one derivative is negative (red) and another is positive (blue). $\partial V^*/\partial\tau$ is indicated by a "*", because it depends on the source of the debt benefits. It is positive if the source of debt benefits is direct. It is ambiguous if the source is the tax shield. Though negative in a wide parameter region, it can be positive, too. † Small region. Less than 2% of the parameter space. †† Minuscule region. Less than 0.001% of the parameter space. Considered effectively unambiguous in the text.

Panel A: Key Comparative Statics on Value and Leverage

		Redeploy-ability η	Reorganization Cost ϕ	Direct Debt Benefits τ
Optimized Firm Value	V^*	0.9,0.2,0.9† 0.9,0.0,0.1	$\downarrow$	*
Debt Face Value, Industry	F^*_{Ind}	0.6,0.0,0.1 0.1,0.2,0.1	$\downarrow$	$\uparrow$
Low-Debt Firm	F^*_L		0.1,0.2,0.1 0.6,0.0,0.1†	0.02,0.02,0.02†† 0.1,0.2,0.1
Debt , Industry	D^*_{Ind}	0.8,0.0,0.1 0.1,0.9,0.1	$\downarrow$	0.9,0.9,0.9 0.1,0.3,0.1
Low-Debt Firm	D^*_L		0.1,0.2,0.1 0.6,0.0,0.1	
Debt / Value, Industry	D^*_{Ind}/V^*	0.8,0.0,0.1 0.1,0.9,0.1	0.1,0.2,0.1 0.9,0.5,0.5	0.1,0.1,0.1 0.1,0.4,0.1
Low-Debt Firm	D^*_L/V^*			

Panel B: Ancillary Comparative Statics

Low Type Credit Spread	$r(F^*_L)$	0.4,0.1,0.4 0.1,0.2,0.1	0.1,0.2,0.1 0.08,0.0,0.02†	0.02,0.02,0.02†† 0.1,0.2,0.1
Asset Price	P^*	$\uparrow$	$\uparrow$	$\downarrow$
Asset Price/Max Value (NPV 0)	P^*/η	0.1,0.2,0.1 0.4,0.3,0.5	$\uparrow$	$\downarrow$
Asset Sales #	Q^*	$\uparrow$	$\uparrow$	0.6,0.0,0.1 0.1,0.2,0.1
Low Type Liquidation Freq.	$\Lambda(F^*_L)/F^*_L$	$\uparrow$	$\uparrow$	$\downarrow$
Low Type Reorganization Cost	$E_v[\phi \cdot (F^*_L - v) \cdot \mathbb{1}_{\Lambda(F^*_L) \le v \le F^*_L}]$	$\downarrow$	0.1,0.2,0.1 0.9,0.0,0.8	0.02,0.02,0.02†† 0.1,0.2,0.1

Table 3.3: Model Implication and Features Comparison

Table 3.3 compares the features and implications of our model with those of the related literature.

	Comparative Statics of Industry Indebtedness Measure			Model Features	
	$\frac{\partial \text{Leverage } D/V}{\partial \text{Debt Benefits}}$	$\frac{\partial \text{Level } D}{\partial \text{Debt Benefits}}$	$\frac{\partial \text{Indebtedness}}{\partial \text{Redeployability}}$	Endogenous Asset Price	Hetero-geneity
Williamson 1988	D/V not derived	Positive[(a)]	Positive[(b)]	No	No
Harris-Raviv 1990	D/V derived, but benefits unexplored[(c)]	Benefits unexplored	Positive[(d)]	No	No
Shleifer-Vishny 1992	D/V not derived	Negative within parameter region. Positive across.[(e)]	Positive[(f)]	Mostly[(g)]	Exogenous[(h)]
Acharya-Vish-wanathan 2011	D/V not derived	Negative for existing firms.[(i)] Positive for new firms.	Redeployability online only. No comparative statics.	Yes	Exogenous
Our Model	Positive when debt benefits τ are small. Negative when large.[(j)]	Deemphasized due to empirical near-un-identifiability. Ambiguous.	Negative when acquisition channel dominates. Positive when liquidation channel dominates.	Yes	Endogenous when assets are indivisible

---

(a) When debt is simpler to implement, more firms choose debt over equity (cf. pg. 579-581).

(b) Equity complexity is necessary for specific hard-to-transfer assets. Firms prefer debt when assets are easy to liquidate (cf. pg. 579-581).

(c) The benefits of debt are that payment/non-payment and audits in default provide signals of asset quality (cf. pg. 329).

(d) Only one firm. Redeployability is a discount in liquidation relative to fundamental value (cf. pg. 340).

(e) Debt helps avoid negative NPV investment and is firm specific. (Own) debt decreases within and increases between equilibrium regions (cf. pg. 1354).

(f) Discussed only informally relative to Williamson 1988 (cf. pg. 1359).

(g) Endogenous decision to sell to insider/outsider. Only when sold to the outsider is the price dependent on debt levels. Otherwise exogenous. (maybe cf. pg. 1353).

(h) The model cannot accommodate identical firms.

(i) They investigate responses to increases in the good asset's expected quality (not relative to bad investment) and show that it decreases debt levels on the intensive margin but increases on the extensive margin.(cf. pg. 120, para. 4).

(j) Measured as net (parameter τ). The value continues to increase even though firms begin to max out leverage.

3.11 Bibliography

Acharya, Viral V. and S. Viswanathan (2011). "Leverage, Moral Hazard, and Liquidity." *The Journal of Finance* 66.1, pp. 99–138.

Acharya, Viral, Sreedhar Bharath, and Anand Srinivasan (2007). "Does industry-wide distress affect defaulted firms? Evidence from creditor recoveries." *Journal of Financial Economics* 85, pp. 787–821.

Allen, Franklin and Douglas Gale (1994). "Limited Market Participation and Volatility of Asset Prices." *American Economic Review* 84.4, pp. 933–955.

Asquith, Paul, Robert Gertner, and David Scharfstein (1994). "Anatomy of financial distress: An examination of junk-bond issuers." *The Quarterly Journal of Economics* 109, pp. 625–658.

Benmelech, Efraim and Nittai K. Bergman (2008). "Liquidation Values and the Credibility of Financial Contract Renegotiation: Evidence from U.S. Airlines." *Quarterly Journal of Economics* 123.4, pp. 1635–1677.

Benmelech, Efraim, Mark J. Garmaise, and Tobias J. Moskowitz (2005). "Do liquidation values affect financial contracts? Evidence from commercial loan contracts and zoning regulation." *Quarterly Journal of Economics* 120.3, pp. 1121–1154.

Bolton, Patrick, Tano Santos, and Jose Scheinkman (2011). "Outside and Inside Liquidity." *The Quarterly Journal of Economics* 126.1, pp. 259–321.

Bris, Arturo, Ivo Welch, and Ning Zhu (2006). "The costs of bankruptcy: Chapter 7 liquidation versus chapter 11 reorganization." *The Journal of Finance* 61 (3), pp. 1253–1303.

Duffie, Darrell (2010). "Presidential Address: Asset Price Dynamics with Slow-Moving Capital." *The Journal of Finance* 65, pp. 1237–1267.

Fries, Steven, Marcus Miller, and William Perraudin (1997). "Debt in industry equilibrium." *Review of Financial Studies* 10, pp. 39–67.

Gale, Douglas and Piero Gottardi (2011). "Bankruptcy, Finance Constraints, and the Value of the Firm." *American Economic Journal: Microeconomics* 3 (2), pp. 1–37.

Gale, Douglas and Piero Gottardi (2015). "Capital Structure, Investment, and Fire Sales." *Review of Financial Studies* 28 (9), pp. 2502–2533.

Gryglewicz, Sebastian (2011). "A theory of corporate financial decisions with liquidity and solvency concerns." *Journal of Financial Economics* 99, pp. 365–384.

Harris, Milton and Artur Raviv (1990). "Capital structure and the informational role of debt." *Journal of Finance* 45, pp. 321–349.

Jensen, Michael and William Meckling (1976). "Theory of the firm: managerial behavior, agency costs and ownership structure." *Journal of Financial Economics* 4, pp. 305–360.

Leary, Mark T. and Michael R. Roberts (2014). "Do Peer Firms Affect Corporate Financial Policy?" *Journal of Finance* 69 (1), pp. 139–178.

Leland, Hayne (1994). "Corporate debt value, bond covenants, and optimal capital structure." *Journal of Finance* 49, pp. 1213–1252.

Leland, Hayne and Klaus Bjerre Toft (1996). "Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads." *Journal of Finance* 51, pp. 987–1019.

Maksimovic, Vojislav and Josef Zechner (1991). "Debt, agency costs, and industry equilibrium." *Journal of Finance* 46, pp. 1619–1643.

Marquez, Robert and M. Denis Yavuz (2013). "Specialization, productivity, and financing constraints." *Review of Financial Studies* 26, pp. 2961–2984.

McDonald, Robert and Donald Siegel (1986). "The Value of Waiting to Invest." *Quarterly Journal of Economics* 101.4, pp. 707–727.

Miller, Merton H. (1977). "Debt and Taxes." *The Journal of Finance* 32.2, pp. 261–275.

Morellec, Erwan and Alexei Zhdanov (2008). "Financing and takeovers." *Journal of Financial Economics* 87.3, pp. 556–581.

Myers, Stewart C. (1977). "Determinants of Corporate Borrowing." *Journal of Financial Economics* 5.2, pp. 147–175.

Opler, Timothy C. and Sheridan Titman (1994). "Financial distress and corporate performance." *Journal of Finance* 49, pp. 1015–1040.

Pulvino, Todd (1998). "Do asset fire sales exist? An empirical investigation of commercial aircraft sale transactions." *Journal of Finance* 53, pp. 939–978.

Rajan, Raghuram and Rodney Ramcharan (2016). "Local financial capacity and asset values: Evidence from bank failures." *Journal of Financial Economics* 120, pp. 229–251.

Robichek, Alexander A. and Stewart C. Myers (1966). "Problems in the Theory of Optimal Capital Structure." *The Journal of Financial and Quantitative Analysis* 1.2, pp. 1–35.

Shleifer, Andrei and Robert Vishny (1992). "Liquidation values and debt capacity: An equilibrium approach." *Journal of Finance* 47, pp. 1343–1366.

Titman, Sheridan (1984). "The effect of capital structure on a firm's liquidation decision." *Journal of Financial Economics* 13, pp. 137–151.

Williamson, Oliver E. (1988). "Corporate finance and corporate governance." *Journal of Finance* 43, pp. 567–592.