Monopolists and Viscous Demand

Roy Radner
Leonard N. Stern School of Business
New York University

Thomas J. Richardson
Bell Laboratories

August 1999

Monopolists and Viscous Demand

Roy Radner
Stern School of Business
New York University

Thomas J. Richardson
Bell Laboratories
Murray Hill, NJ 07974

June 23, 1997

ABSTRACT

We characterize the optimal dynamic price policy of a monopolist who faces "viscous" demand for its services. Demand is *viscous* if it adjusts relatively slowly to price changes. We show that with the optimal policy the monopolist stops short of achieving 100% market penetration, even when all of the consumers have the same long-run willingness to pay for the service. Furthermore, for certain parameter values in the model, the price policy requires rapid oscillations of the price path.

Monopolists and Viscous Demand

1. Introduction and Summary

1.1. Introductory Remarks

"Viscous demand" refers to a phenomenon common to many markets, in which demand adjusts relatively slowly to changes in prices and quality. There are many reasons for the viscosity of demand, depending on the type of good or service being sold, and on the organization of the market. In this paper we focus on the case of a service, such as a subscription to a magazine, newspaper, or long-distance carrier, that does not absorb a major part of the customer's budget. We also focus on responses to price changes, rather than changes in marketing or quality, or the introduction of innovations. In a sense, in this case the viscosity of demand is due to a certain type of bounded rationality and myopia on the part of the consumers. Although this case is somewhat special, we hope that the approach we introduce here will be helpful in other models. In particular, one of the parameters of our model can be interpreted as the effectiveness of marketing in inducing consumers to change behavior. (For a more general introduction to the concept and causes of demand viscosity, see Radner, 1997.)

The presence of demand viscosity creates a temporary monopoly power for a supplier of the service. In such a case, the time path of a firm's prices acquires added significance, and the problem of optimal pricing becomes significantly more complex, compared with the case of instantaneous demand response. In particular, as one might expect, the sequential-game-theoretic problems inherent in a thorough and rigorous treatment of oligopolistic markets are an order of magnitude more difficult than in the instantaneous-demand case. In fact, there is very little previous theoretical literature on this topic, and – in our opinion – the models described in this literature have not been satisfactory. (For a review of this literature, and an analysis of a model of duopoly, see Radner, 1997.)

For these reasons, we focus in this paper on the case of a monopolist selling a service that can be bought by each consumer at a fixed rate per unit of time, or not at all. (The precise model is described below.) In spite of the simplicity of the situation described by the model, and the relatively small number of parameters, we shall see that the monopolist's optimal pricing strategy can be radically different for different parameter values. Roughly speaking, we can summarize our results as follows. In one region of the parameter space, call it A, the

1

monopolist has a "target" market penetration, *which is less than 100 percent.* If the initial penetration is below the target, then the firm charges a minimum price (say zero) until the market penetration reaches the target, and then switches to a price that stabilizes the market penetration at the target level. If the initial market penetration is above the target, then a price is charged that stabilizes the market penetration at the initial level. In particular, the parameters will be in the region A if the firm's discount rate is not too high compared to the inverse of the demand viscosity, and the price is bounded above by a sufficiently low bound.

In the complementary region of the parameter space, call it B, the firm's price oscillates rapidly between its lower and upper bounds, maintaining its average market penetration at a low level. In fact, *strictly speaking there is no optimal policy in the usual sense; we may say that in the "optimal" policy the price oscillates infinitely fast between its lower and upper bounds!* (We provide a precise meaning to this statement.) Such a situation is hardly realistic, and provokes a reconsideration of the behavioral assumptions of the model when it has these parameter values. In particular, if prices are oscillating very quickly, one would not expect consumers to react so myopically as they do in the first postulated model. For example, one might expect (boundedly rational) consumers to forecast prices in some "adaptive" manner, e.g., with a moving average of past prices. We formulate this idea in a second model. Although a complete characterization of the monopolist's optimal policy in the face of consumers with adaptive expectations is not known, we show that (1) an optimal policy exists, (2) the maximum profit in the "adaptive expectations" model approaches that in the "myopic behavior" model as the speed of adaptation increases without bound, and (3) the optimal adaptive-expectations price policy is nearly optimal for the first model if the speed of adaptation is large enough. This last point implies that in the adaptive-expectations model the optimal price may oscillate, but at a finite frequency.

In the remainder of this section, we summarize the assumptions and results for the two models. In Section 2 we present a detailed statement of the models and the main results, and show how the first, or "main," model is a limiting case of the second, or "adaptive-expectations," model. In Section 3 we derive the characterization of the value function and optimal price policy for the main model.

2

1.2. Summary of Results for the Main Model

A monopolist sells a service. The population of consumers is a continuum, with total mass one. Time is continuous. At each instant of time each consumer is purchasing the service at a rate of either 0 or 1; in the latter case the consumer is called a *customer*. (Note that here and henceforth 'customer' refers to a customer of the service.)

For each time $t \geq 0$, let $Q(t)$ denote the mass of customers, and $P(t)$ denote the price of the service per unit time. The mass of customers evolves according to the differential equation

$$\dot{Q}(t) = f(P(t), \ Q(t)) \ , \tag{1.1}$$

where

$$f(p,q) = \begin{cases} \lambda(w-p)(1-q) \ , & p \leq w \ , \\ -\lambda(p-w)q \ , & p \geq w \ , \end{cases} \tag{1.2}$$

and λ and w are strictly positive constants. (We will use upper case to denote dynamic variables and lower case for static and 'dummy' variables.) The essential character of (1.1)–(1.2) is that customers do not respond instantaneously to a change in the price. Thus suppose that in some time interval the price is held constant at the level p. If $p < w$ then the firm will *gain* customers at a rate proportional to the difference, $(w - p)$, and to the mass of remaining *noncustomers*, $(1 - q)$. On the other hand, if $p > q$, then the firm will *lose* customers at a rate proportional to $(p - w)$ and to the mass of remaining *customers*. The constant of proportionality, lambda, is the inverse of the viscosity. We may interpret w as the consumers' "long-run" or "static" willingness to pay for the service. (Note that w is here the same for all consumers.)

The monopolist's instantaneous cost per unit of time at t is

$$k + cQ(t) \ ,$$

where $k, c \geq 0$. The magnitude of k (the fixed cost) does not affect the optimal pricing policy, although it does influence whether the monopolist's profit is positive or not. Hence we shall take $k = 0$.

The monopolist's total discounted profit is therefore

$$V = \int_0^\infty e^{-\rho t}[P(t) - c]Q(t)dt \ , \tag{1.3}$$

where $\rho > 0$ is the discount rate. Given the initial mass of customers, $Q(0)$, the monopolist wants to choose the price path $P(t)$ to maximize (1.3). For reasons that will become apparent

below, we make the following assumptions:

$$0 \leq P(t) \leq \bar{p} , \qquad (1.4a)$$

$$w, c < \bar{p} . \qquad (1.4b)$$

(In place of (1.4a), one could modify (1.2) so that $f(p, q) = -\infty$ for $p > \bar{p}$.)

We will present a complete solution to this problem in the next section. Since the problem is time invariant we need only consider Markov policies, i.e., $P(t) = \phi(Q(t))$, by Blackwell's Theorem. There are essentially two candidate optimal strategies. The first is

$$\phi^1(q) = \begin{cases} 0, & 0 \leq q < a , \\ \bar{p}, & a < q \leq 1 , \end{cases} \qquad (1.5)$$

where

$$a = \begin{cases} \dfrac{w}{2h} \dfrac{\sqrt{\rho^2 + 4\lambda(\bar{p} - w)h} + \rho - 2h}{\bar{p} - 2w} , & \bar{p} \neq 2w , \\[4mm] \dfrac{\lambda(w - \frac{c}{2})}{(\rho + 2\lambda w)} , & \bar{p} = 2w , \end{cases} \qquad (1.6)$$

and $h = \rho + \lambda w + \frac{c}{\bar{p}}\lambda(\bar{p} - 2w)$, and at $q = a$ the policy is to oscillate between $\bar{p}$ and 0 such that $q = a$ is a stationary point. (This will be made more precise in the next section.) The second candidate is

$$\phi^2(q) = \begin{cases} 0, & 0 \leq q < \sigma , \\ w, & \sigma \leq q \leq 1 , \end{cases} \qquad (1.7)$$

$$\sigma = \frac{\lambda(w - c)}{\rho + \lambda(w - c)} . \qquad (1.8)$$

Each of these strategies has associated to it a return, i.e., a discounted profit, $V^1(Q_0), V^2(Q_0)$ respectively. This is the main result of this paper: it is optimal to choose whichever of the two strategies gives the largest return. Depending on the parameters, either one strategy dominates the other for all Q_0, or there is a single value $Q_0 = x$ such that for ϕ^1 dominates for $Q_0 > x$ and ϕ_2 dominates for $Q_0 < x$. Furthermore, if $Q(0) \gtreqless x$ then $Q(t) \gtreqless x$ for all t under the optimal policy, so there is no switching between strategies.

Remark 1. Call $Q(t)$ the (*market*) *penetration*, and call σ the *target* penetration under policy ϕ^2. If the initial penetration is strictly less than the target, then, under the policy ϕ^2, the penetration will increase monotonically to the target, reaching it in finite time. On the other hand, any penetration greater than or equal to the target is a steady state. These conclusions hold even if the target does not satisfy (1.8), i.e., even if it is not optimal.

Remark 2. The optimal target, in (1.8), is decreasing in the marginal cost, c, so that, under

4

policy ϕ^2, the monopolist increases his steady-state penetration if he reduces his marginal cost.

Remark 3. Under policy ϕ^1 the target penetration is a. If $Q(0)$ is larger than a then it is optimal to charge a high price ($\bar{p}$), loosing customers until, in finite time, the penetration reaches a. Here we clearly see the optimal policy taking advantage of the viscosity of the customer base.

Remark 4. As $\bar{p} \to \infty$ the target penetration a in (1.5) goes to 0 like $\bar{p}^{(-\frac{1}{2})}$. Furthermore, policy ϕ^1 dominates policy ϕ^2 for $\bar{p}$ large enough. Here we see a degeneracy in the model. In reality, too large a price would cause a mass exodus of the customer base. To keep the model tractable we let $\bar{p}$ be the price above which all customers leave instantaneously.

Remark 5. It is interesting to note that in the event $V^1(x) = V^2(x)$ for $x \in (0,1)$ then if $Q(0) > x$, i.e. if the customer base is relatively large, then the optimal strategy is to hold or increase the customer base. If the customer base is small, i.e. $Q(0) < x$, then it is optimal to take advantage of its viscosity by overcharging.

1.3. Adaptive Expectations

There is perhaps something unnatural about the solution described above. This unnaturalness lies in the implied behavior of the customer when $Q(t) = a$ and ϕ^1 is optimal. From the customer's perspective prices apparently fluctuate infinitely fast. Consequently, customers flow to and from the service, the net flow being equal to zero. In reality customers will not respond this way. Nevertheless, the 'oscillating' fixed point has significance. If some smoothing mechanism is introduced into the model then the 'infinite' oscillation may disappear and yet the basic character of the oscillation will remain.

We will consider a particular modification to our model. This modification is referred to as "adaptive expectation" and the resulting optimization problem will be referred to as the γ-AR problem. Here γ is a new parameter. We will use ∞-AR to refer to the problem described earlier in which adaptive expectation is absent.

The modification is the following, we put

$$\dot{Q}(t) = f(\tilde{P}(t),\ Q(t))$$

in place of (2.1) where $\tilde{P}(t)$ solves

$$\dot{\tilde{P}}(t) = \gamma(P(t) - \tilde{P}(t))\ ,\quad \tilde{P}(0) = \tilde{P}_0\ .$$

Here, as before, $P(t) \in \psi$ is the price we set and we assume $\tilde{P}_0 \in [0,\ \bar{p}]$.

5

There are two possible interpretations for this modification. The first is that customers do not track prices exactly but have some estimate $\tilde{P}$ based on the recent past. The second interpretation is that some regulatory mechanism dampens price fluctuations and that this mechanism can be modeled as adaptive expectation. In the later case we should optimize $V(Q_0, \tilde{P})$, rather than $V(Q_0, P)$ since $\tilde{P}$ is the price actually paid. However, for large γ we will see that the two are nearly the same (see Eq. (2.2)).

We are not able to solve the γ-AR model explicitly. However, we can and will prove that the γ-AR model admits a real valued optimal solution. That is, unlike the ∞-AR model, infinitely fast oscillations are never required. We will also prove that the γ-AR model approximates the ∞-AR model in the sense that $\tilde{P}_\gamma(t)$ is nearly optimal for ∞-AR if γ is large enough. This will imply the optimal pricing in the γ-AR can require oscillation (at a finite rate) between high ($\bar{p}$) and low (0) prices. Periods of low prices, i.e. sales, draw in customers. Subsequently prices are raised; customers leave, but not instantaneously and profits are reaped. This behavior can be seen in the 'real world.'

2. Formulation and Statement of Main Results

We will now present a mathematically precise statement of our ∞-AR problem. Let ψ denote the set of measurable functions from $[0, \infty)$ to $[0, \bar{p}]$. Given $P \in \psi$ and $Q_0 \in [0, 1]$ let $Q(t)$ be the solution to

$$
\begin{aligned}
\dot{Q}(t) &= f(P(t), Q(t)), \\
Q(0) &= Q_0,
\end{aligned}
\tag{2.1}
$$

where

$$
f(p, q) := \begin{cases} \lambda(w - p)(1 - q), & p \leq w, \\ -\lambda(p - w)q, & p > w. \end{cases}
$$

Define

$$
V(Q_0, P) := \int_0^\infty e^{-\rho t}(P(t) - c)Q(t)dt
$$

and

$$
V^{opt}(Q_0) := \sup_{P \in \psi} V(Q_0, P) .
$$

Our goal is to find V^{opt} and, if possible, find a Markov policy $P(t) = \phi(Q(t))$ which achieves V^{opt}. By this we mean a measurable function $\phi : [0, 1] \to [0, \bar{p}]$ such that if $Q(t)$ solves (2.1) with $P(t)$ formally set to $\phi(Q(t))$ then $V(Q_0, \phi(Q(t))) = V^{opt}(Q_0)$. It will turn out that this is

6

not possible in general and it is necessary to broaden the class of admissible controls to admit measure valued controls [Ga].

Typically in problems of this type one broadens the class of admissible controls in order to obtain existence results and then proves that a 'regular' optimal control can be found in the original class. This is usually achieved by virtue of certain convexity properties of the functional to be minimized. Our functional, it turns out, lacks these properties so that $P(t)$ must necessarily be replaced by a measure valued control μ_t. We can then find an optimal Markov policy of the form $\mu_t = \phi(Q(t))$.

2.1. The Weak Formulation: Measure Valued Pricing

Let μ_t be a family of probability measures on $[0, \bar{p}]$ depending on the parameter $t \in [0, \infty)$. Let $g(t, p)$ be a continuous function on $[0, \infty) \times [0, \bar{p}]$. We define the function

$$h(t) := E_{\mu_t}(g(t, P_t))$$

where P_t denotes a random variable with law μ_t. If $h(t)$ is Lebesgue measurable for an arbitrary continuous g then we say μ_t, $t \in [0, \infty)$ is weakly measurable with respect to t. Such an object is said to be a *generalized control*. We denote the set of generalized controls by Θ.

The weak formulation of the optimization problem is the following. Maximize

$$\tilde{V}(Q_0, \mu_t) := \int_0^\infty e^{-\rho t}(E_{\mu_t}(P_t) - c)\, Q(t) dt$$

where $Q(t)$ solves

$$\dot{Q}(t) = E_{\mu_t}(f(P_t, Q(t))),$$

$$Q(0) = Q_0.$$

Define

$$\tilde{V}^{opt}(Q_0) = \sup_{\mu_t \in \Theta} \tilde{V}(Q_0, \mu_t).$$

There are a few fundamental results relating the weak formulation to the original formulation. The first is

$$\tilde{V}^{opt}(Q_0) = V^{opt}(Q_0).$$

The inequality $\tilde{V}^{opt}(Q_0) \geq V^{opt}(Q_0)$ follows from the fact that $\delta_{P(t)} \in \Theta$ for $P \in \psi$. Here δ_p is the Dirac delta function at $p \in [0, \bar{p}]$. The inequality $\tilde{V}^{opt}(Q_0) \leq V^{opt}(Q_0)$ can be proved by finding a sequence $P^n(t)$ converging "weakly" (in the sense of generalized controls) to μ_t and showing that $V(Q_0, P^n(t)) \to \tilde{V}(Q_0, \mu_t)$. The advantage the weak formulation has over the

7

original formulation arises from the fact that Θ is a convex space whereas ψ is not. General lower semicontinuity and compactness properties of generalized controls [Ga] allow us to assert

Theorem 2.1. *For any $Q_0 \in [0,1]$ there exists an optimal generalized control μ_t^{opt}, i.e.,*

$$V^{opt}(Q_0) = \tilde{V}(Q_0, \mu_t^{opt}) \ .$$

It turns out that our problem has a piecewise affine structure which allows us to restrict to a small subclass of Θ. Given a probability measure μ on $[0, \bar{p}]$ there exist constants $m_0 \geq 0, m_{\bar{p}} \geq 0$ satisfying $m_0 + m_{\bar{p}} \leq 1$ such that, for any continuous function g on $[0, \bar{p}]$ which is affine on $[0, w]$ and on $[w, \bar{p}]$, we have

$$E_\mu(g) = m_0 \ g(0) + (1 - m_0 - m_{\bar{p}}) \ g(w) + m_{\bar{p}} \ g(\bar{p}) \ .$$

Now, $f(p, q)$ and $p - c$ are piecewise affine in this sense (in p). Thus, any generalized control μ_t can be replaced by a control

$$\tilde{\mu}_t = m_0(t) \ \delta_0 + (1 - m_0(t) - m_{\bar{p}}(t)) \ \delta_w + m_{\bar{p}}(t) \ \delta_{\bar{p}}$$

without altering $Q(t)$ or $\tilde{V}(Q_0, \cdot)$. (Strictly speaking we should verify $\tilde{\mu}_t \in \Theta$: It can easily be verified that $m_0(t)$ and $m_{\bar{p}}(t)$ are measurable.) Define

$$\Omega = \{(m_0, m_{\bar{p}}) : m_0 \geq 0, \ m_{\bar{p}} \geq 0, \ m_0 + m_{\bar{p}} \leq 1\} \ .$$

We now have

Theorem 2.2. *For any $Q_0 \in [0,1]$ there exist a measurable function $(m_0, m_{\bar{p}}) : [0, \infty) \to \Omega$ satisfying*

$$V^{opt}(Q_0) = \tilde{V}(Q_0, \tilde{\mu}_t^{opt}) ,$$

where $\tilde{\mu}_t^{opt} = m_0(t) \ \delta_0 + (1 - m_0(t) - m_{\bar{p}}(t)) \ \delta_w + m_{\bar{p}}(t) \ \delta_{\bar{p}}$.

We will look for a Markov optimal policy, that is $\phi : [0, 1] \to \Omega$ such that setting $(m_0(t), m_{\bar{p}}(t)) = \phi(Q(t))$ yields an optimal control. By Blackwell's theorem this is sufficient.

2.2. Statement of Main Result: Optimal Policy

We are now ready to present the optimal value function V^{opt} and an optimal policy ϕ. Define

$$D(x) = f(0, x) - f(\bar{p}, x)$$

8

and the quadratic function

$$Q(x) := (\lambda w)^2(1-x) - (\lambda w + \rho)xD(x) - \frac{c}{\bar{p}}D^2(x).$$

Lemma 2.3. *Q has exactly one root a in $[0,1]$ and $sgn(Q(x) - a) = -sgn(x - a)$ on $[0,1]$.*

Proof. Since $Q(0) = (\lambda w)^2(1 - \frac{c}{\bar{p}}) > 0$ and $Q(1) = -(\lambda w + \rho)(\lambda(\bar{p} - w)) - \frac{c}{\bar{p}}D^2(1) < 0$, we see that Q has exactly one root in $[0,1]$. $\square$

Remark. We have given a formula for a in (1.6).

We now define two possible value functions V^1 and V^2. The function V^1 is the value associated with the following policy

$$\phi^1(q) := \begin{cases} m_0 = 1, & m_{\bar{p}} = 0, & q < a\,, \\ m_0 = 1 - m(a), & m_{\bar{p}} = m(a), & q = a\,, \\ m_0 = 0, & m_{\bar{p}} = 1, & q > a\,. \end{cases}$$

and the function V^2 is associated with the policy

$$\phi^2(q) := \begin{cases} m_0 = 1, & m_{\bar{p}} = 0, & q < \sigma\,, \\ m_0 = 0, & m_{\bar{p}} = 0, & q \geq \sigma\,, \end{cases}$$

where

$$\sigma = \frac{\lambda(w - c)}{\lambda(w - c) + \rho}\,,$$

Closed form expressions for V^1 and V^2 can be easily obtained. Let ϕ^{opt} be a policy where $\phi^{opt} = \phi^1$ when $V^1 > V^2$ and $\phi^{opt} = \phi^2$ when $V^2 > V^1$ and ϕ^{opt} is either ϕ^1 or ϕ^2 when $V^1 = V^2$. The following is the main result to be proved.

Theorem 2.4. [Optimality] *The policy ϕ^{opt} is optimal and*

$$V^{opt} = \max(V^1, V^2)\,.$$

Furthermore either $V^{opt} = V^1$ or $V^{opt} = V^2$ or there exists $x \in (a, 1)$ such that $V^{opt} = V^1$ on $[0, x]$ and $V^{opt} = V^2$ on $[x, 1]$.

Remark. The function V^{opt} is C^2 except at x where it fails to be differentiable.

2.3. Regularization Via the γ-AR Model

Our primary purpose in considering the γ-AR model is to lend credence to the limiting case where we have a complete solution. With $\gamma < \infty$ the solution is "regularized" while retaining its basic structure.

9

2.4. Existence of Optimal Control for the γ-AR Model

Establishing existence of optimal control functions for the γ-AR model is straightforward. Let μ_t be the "weak" limit of a minimizing sequence from ψ. It follows easily that μ_t is a weak minimizer. However, $P(t) := E(\mu_t)$ is obviously equivalent to μ_t with respect to the γ-AR model. This establishes existence.

An alternative proof that avoids generalized controls proceeds as follows. Since $\tilde{P}_\gamma(t)$ is Lipshitz (with constant $\gamma \bar{p}$) we can extract a minimizing sequence which converges uniformly on $[0, T]$ for any $T < \infty$ to some Lipshitz function $\tilde{P}_\gamma^*(t)$. We then define $P(t) = \tilde{P}_\gamma^*(t) + \frac{1}{\gamma} \frac{d}{dt} \tilde{P}_\gamma^*(t)$ and verify $P(t) \in [0, \bar{p}]$ for almost all t. It follows that $P(t)$ is optimal.

2.5. Approximation Results

It is fairly easy to prove that V_γ^{opt} approximates V_∞^{opt} for large γ.

Theorem 2.5. *For any $\tilde{P}(0) \in [0, \bar{p}]$ we have*

$$\lim_{\gamma \to \infty} V_\gamma^{\text{opt}}(Q_0, \tilde{P}(0)) = V_\infty^{\text{opt}}(Q_0) \ .$$

Proof. If $P(t)$ is a C^∞ function then it is easy to prove

$$\lim_{\gamma \to \infty} V_\gamma(C_0, \tilde{P}(0), P(t)) = V_\infty(Q_0, P(t)) \ .$$

Since there exist $P(t) \in C^\infty$ such that $V_\infty^{\text{opt}}(Q_0) \leq V_\infty(Q_0, P(t)) + \epsilon$, where $\epsilon > 0$ is arbitrary, we have $\lim_{\gamma \to \infty} V_\gamma^{\text{opt}}(Q_0, \tilde{P}(0)) \geq V_\infty^{\text{opt}}(Q_0)$. The opposite inequality is trivial since $V_\gamma(Q_0, P(\tilde{0}), P(t)) = V_\infty(Q_0, \tilde{P}(t))$. $\quad \square$

Now, suppose $\{P_\gamma(t)\}_{\gamma=1,2\ldots}$ is a sequence of optimal pricing strategies for the initial condition $Q_0, \tilde{P}(0)$. By weak compactness of generalized controls there exists μ_t, a generalized control which is the "weak" limit of some subsequence of $\{\tilde{P}_\gamma(t)\}_{\gamma=1,2,\ldots}$. It follows that μ_t is optimal for the non-adaptive problem. To see this: By integrating by parts we have $|\int_0^\infty e^{-\rho t} Q_\gamma(t) \dot{P}_\gamma(t) dt| \leq C$ for some constant C. Hence,

$$\left| \int_0^\infty e^{-\rho t} Q_\gamma(t) (\tilde{P}_\gamma(t) - P_\gamma(t)) dt \right| \leq \frac{C}{\gamma} \ . \tag{2.2}$$

Applying this inequality we obtain

$$\lim_{\gamma \to \infty} V_\gamma^{\text{opt}}(Q_0, \tilde{P}(0)) = \lim_{\gamma \to \infty} V_\infty(Q_0, \tilde{P}(t)) = V_\infty(Q_0, \mu_t) \ .$$

We conclude that μ_t is optimal from Theorem 2.5. Thus, we have proved

Theorem 2.6. *Any "weak" cluster point (in the sense of generalized controls) of optimal $\{\tilde{P}_\gamma(t)\}$ as $\gamma \to \infty$ is optimal for the ∞-AR problem.*

10

3. Main Result

In this section we prove our solution to the ∞-AR problem. Herein we shall therefore dispense with the γ-∞ distinction.

3.1. The Value Function: Viscosity Solutions

For $\xi, q, p \in \mathbb{R}$ define

$$G(\xi, q, p) := \xi f(p, q) + (p - c)q$$

and

$$F(\xi, q) := \sup_{p \in [0, \bar{p}]} G(\xi, q, p) .$$

For future reference we also define

$$\tilde{G}(\xi, q, m_0, m_{\bar{p}}) := \xi(m_0 f(0, q) + m_{\bar{p}} f(\bar{p}, q)) + ((1 - m_0)w + m_{\bar{p}}(\bar{p} - w) - c)q .$$

Further, for $u \in \mathbb{R}$, we define the 'Hamiltonian'

$$H(q, u, \xi) := \rho u - F(\xi, q) .$$

Formally, the value function V^{opt} satisfies the Hamilton-Jacobi-Bellman equation,

$$H(q, \mathcal{U}, \mathcal{U}_q) = 0 \qquad \text{(HJB)}$$

where $\mathcal{U}$ is a function on $[0, 1]$. This equation does not have a solution in the classical sense, in general. Fortunately, in the last 15 years the theory of "viscosity solutions" has been developed. We will not define the concept of a viscosity solution here and refer the interested reader to [FS] [Li] [LS] [CHL].[1] The main point is the following

Theorem 3.1. V^{opt} *is the unique viscosity solution to HJB.*

Any classical, i.e. C^1, solution to HJB is a viscosity solution, hence must be V^{opt}. For our problem V^{opt} will not be C^1 in general and the classical theory of such equations is insufficient. It turns out that V^{opt} is locally a classical solution except possibly at the point x. Only for this point are we required to verify non-classical 'viscosity' conditions of the solution. It turns out that a continuous piecewise C^1 function which satisfies HJB in the classical sense on those (closed) intervals determined by the finite number of points where the function is not locally C^1 is automatically the viscosity solution (see [Li] or [LS]).

---

[1]The references use the formalism of minimization rather than maximization to define the value function. Hence, appropriate sign changes are required to correctly interpret them.

3.2. Proof of Main Result

Define the function

$$H^1(q) \ := \ \frac{m(q)\,\bar{p} - c}{\rho}\, q\,,$$

$$\text{where}\quad m(q) \ := \ \frac{f(0,q)}{f(0,q) - f(\bar{p}, q)}\,.$$

If we choose $m_0 = 1 - m(q)$ and $m_{\bar{p}} = m(q)$ then q is a stationary point under $(m_0, m_{\bar{p}})$, and the corresponding value at q is $H^1(q)$, i.e.,

$$\int_0^\infty e^{-\rho t}((1 - m_0)w + m_{\bar{p}}(\bar{p} - w) - c)q \ dt = H^1(q)\,.$$

Let us also define

$$H^2(q) := \frac{(w - c)q}{\rho}\,.$$

This is the value under the stationary policy $m_0 = 0$, $m_{\bar{p}} = 0$. Note that $V^1(a) = H^1(a)$ and $V^2(q) = H^2(q)$ for $q \in [\sigma, 1]$.

We will consider the policies

$$\phi'_b := \begin{cases} m_0 = 1, & m_{\bar{p}} = 0, & q < b\,, \\[2mm] m_0 = 1 - m(b), & m_{\bar{p}} = m(b), & q = b\,, \\[2mm] m_0 = 0, & m_{\bar{p}} = 1, & q > b\,, \end{cases}$$

and the value function $V(q, \phi'_b)$. Note that $V(b, \phi'_b) = H^1(b)$ and $V^1 = V(\phi'_a)$.

Lemma 3.2. *For $b \neq a$ we have $V(\phi_b) < V^1$ and, furthermore, $V^1 \geq H^1$ with equality holding only at a.*

Proof. We claim

$$sgn\left(\frac{\partial}{\partial b}V(q, \phi'_b)\right) = sgn(a - b)\,. \tag{3.1}$$

The claim proves the lemma since

$$V^1(q) - H^1(q) = \int_q^a \left(\frac{\partial}{\partial b}V(q, \phi'_b)\right) db$$

and

$$V^1(q) - V(q, \phi'_{b'}) = \int_{b'}^a \left(\frac{\partial}{\partial b}V(q, \phi'_b)\right) db\,.$$

We now prove the claim. In general we have

$$V(q, \phi'_b) = e^{-\rho T(q)}H^1(b) + \int_0^{T(q)} e^{-\rho t}(1_{q > b}\bar{p} - c)Q(t)dt$$

12

where $T(q) = \min\{t : Q(t) = b\}$. If $q \geq b$ then $e^{-\lambda(\bar{p}-w)T(q)}q = b$ and if $q \leq b$ then $e^{-\lambda w T(q)}(1 - q) = 1 - b$. Assume for convenience that $q \neq b$. We obtain

$$\frac{\partial}{\partial b}V(q, \phi'_b) = e^{-\rho T(q)}\frac{\partial}{\partial b}H^1(b) + e^{-\rho T(q)}(-\rho H^1(b) + (1_{q>b}\bar{p} - c)b)\frac{\partial}{\partial b}T(q) \ .$$

Now, $\frac{\partial}{\partial b}T(q) = \frac{1}{f(0,b)}$ or $\frac{\partial}{\partial b}T(q) = \frac{1}{f(\bar{p},b)}$ according to $q < b$ or $q > b$. By calculating

$$\frac{\partial}{\partial b}H^1(b) = \frac{\lambda w \bar{p} f(0,b) - \lambda w b D - cD^2}{\rho D^2}$$

and noting that $\frac{1}{f(\bar{p},b)}(-\rho H^1(b) + (\bar{p} - c)b) = \frac{1}{f(0,b)}(-\rho H^1(b) - cb) = \frac{-\bar{p}b}{D(b)}$, we obtain

$$\frac{\partial}{\partial b}V(q, \phi'_b) = \frac{\bar{p}e^{-\rho T(q)}}{\rho D^2(b)}\mathcal{Q}(b) \ .$$

Since this formula is continuous at $q = b$ it holds there also. By Lemma 2.3 equation (3.1) holds. $\square$

We now define the policies

$$\phi''_r := \begin{cases} m_0 = 1, & m_{\bar{p}} = 0 & q < r \\ m_0 = 0, & m_{\bar{p}} = 0 & q \geq r \ . \end{cases}$$

Lemma 3.3. *For all $r \in [0,1]$ we have $V(\phi''_r) \leq V^2$. Furthermore $V^2 \geq H^2$ with equality holding only on $[\sigma, 1]$.*

Proof. For $q \geq \max\{\sigma, r\}$ we have $V(q, \phi''_r) = V^2(q)$. For $q < \max\{\sigma, r\}$ we have $V^2(q) - V(q, \phi''_r) = \int_r^\sigma \left(\frac{\partial}{\partial s}V(q, \phi''_s)\right) ds$. If $q < r$ then $V(q, \phi_r) = e^{-\rho T}H^2(r) + \int_0^T e^{-\rho t}(-c)Q(t)dt$ where $Q(t)$ solves $\dot{Q}(t) = \lambda w(1 - Q(t))$, $Q(0) = q$, and T is given by $Q(T) = r$. Hence, in this case we have

$$\frac{\partial}{\partial r}V(q, \phi_r) = -e^{-\rho T}\frac{\partial}{\partial r}H^2(r) + e^{-\rho T}(-\rho H^2(r) - cr)\frac{\partial}{\partial r}T$$

$$= e^{-\rho T}\frac{\rho + \lambda(w-c)}{\lambda \rho(1-r)}(\sigma - r) \ .$$

For $q \geq r$ we have $\frac{\partial}{\partial r}V(q, \phi''_r) = 0$, hence $V(\phi''_r) \leq V^2$. To prove the last statement we note $H^2(q) = V(q, \phi''_q)$. $\square$

Lemma 3.4. *V^1 satisfies HJB at q if and only if $V^1(q) \geq H^2(q)$.*

Proof. Let

$$F^1(q, m_0, m_{\bar{p}}) := \tilde{G}(V_q^1(q), q, m_0, m_{\bar{p}}) \ .$$

For V^1 to satisfy HJB on $[0, a)$ it is sufficient to show that $\frac{\partial}{\partial m_0}F^1(q, m_0, m_{\bar{p}}) \geq 0$ and $\frac{\partial}{\partial m_0}F^1(q, m_0, m_{\bar{p}}) \geq \frac{\partial}{\partial m_{\bar{p}}}F^1(q, m_0, m_{\bar{p}})$. The first condition reduces to $V_q^1(q)f(0,q) \geq wq$. Since, on $[0, a]$ we have $\rho V^1(q) = V_q^1(q)f(0,q) - cq$ the condition further reduces to $\rho V^1(q) \geq (w-c)q$, or

13

$V^1(q) \geq H^2(q)$. Similarly, the second condition reduces to $V_q^1(q)(f(0,q) - f(\bar{p}, q)) \geq \bar{p}q$ which reduces to $V^1(q) \geq H^1(q)$. Since the second condition is satisfied, according to Lemma 3.2, we see that $V^1(q) \geq H^2(q)$ is necessary and sufficient to guarantee HJB.

For $q = a$ the same argument holds but we require $\frac{\partial}{\partial m_0} F^1(q, m_0, m_{\bar{p}}) = \frac{\partial}{\partial m_{\bar{p}}} F^1(q, m_0, m_{\bar{p}})$, or $V^1(a) = H^1(a)$ which holds by Lemma 3.2.

For $q \in (a, 1)$ we require $\frac{\partial}{\partial m_{\bar{p}}} F^1(q, m_0, m_{\bar{p}}) \geq 0$ and $\frac{\partial}{\partial m_{\bar{p}}} F^1(q, m_0, m_{\bar{p}}) \geq \frac{\partial}{\partial m_0} F^1(q, m_0, m_{\bar{p}})$. Note that here V^1 satisfies $\rho V^1(q) = V_q^1(q)f(\bar{p}, q) + (\bar{p} - c)q$. A calculation shows that the two conditions reduce to the same conditions as those above, namely $V^1(q) \geq H^2(q)$ and $V^1(q) \geq H^1(q)$ respectively. By Lemma 3.2, the proof is complete. $\square$

Lemma 3.5. V^2 satisfies HJB at $q \in [0, \sigma)$ if and only if $V^2(q) \geq H^1(q)$.

Proof. Define

$$F^2(q, m_0, m_{\bar{p}}) := \tilde{G}(V_q^2(q), q, m_0, m_{\bar{p}}) .$$

For V^2 to satisfy HJB at $q \in [0, \sigma)$ it is necessary and sufficient that $\frac{\partial}{\partial m_0} F^2(q, m_0, m_{\bar{p}}) \geq 0$ and $\frac{\partial}{\partial m_0} F^2(q, m_0, m_{\bar{p}}) \geq \frac{\partial}{\partial m_{\bar{p}}} F^2(q, m_0, m_{\bar{p}})$. Since V^2 satisfies $\rho V^2(q) = V_q^2(q)f(0,q) - cq$ here these conditions reduce to $V^2(q) \geq H^2(q)$ and $V^2(q) \geq H^1(q)$ respectively exactly as in Lemma 3.4. By Lemma 3.3 the first condition holds in general. $\square$

Lemma 3.6. V^2 satisfies HJB on $[\sigma, 1]$ if and only if $\sigma \geq \frac{1}{2}$.

Proof. Here we require $\frac{\partial}{\partial m_0} F^2 \leq 0$ and $\frac{\partial}{\partial m_{\bar{p}}} F^2 \leq 0$ where F^2 is as in the proof of Lemma 3.5. Since $V^2(q) = H^2(r)$ these conditions easily reduce to $q \geq \sigma$ and $\sigma \geq \frac{1}{2}$ respectively. The first condition is given, thus, the proof is complete. $\square$

Lemma 3.7. If $\sigma \leq \frac{1}{2}$ then $V^1 \geq V^2$ with strict inequality unless $a = \sigma = \frac{1}{2}$.

Proof. Assume $\sigma \leq \frac{1}{2}$ and consider the policies (assume $r \geq \sigma$)

$$\phi_r'''(q) := \begin{cases} m_0 = 1, & m_{\bar{p}} = 0 \quad q < \sigma \\ m_0 = 0, & m_{\bar{p}} = 0 \quad q \in [\sigma, r] \\ m_0 = 0, & m_{\bar{p}} = 1 \quad q > r . \end{cases}$$

For $q \in [0, \sigma]$ we have $V(q, \phi_r''') = V^2(q)$. For $q > \sigma$ we have $V(q, \phi_\sigma''') - V^2(q) = -\int_\sigma^q \left(\frac{\partial}{\partial r} V(q, \phi_r''') \right) dr$. Since, here, $V(q, \phi_r''') = e^{-\rho T} H^2(r) + \int_0^T e^{-\rho t}(\bar{p} - c)Q(t)dt$ where $Q(t)$ solves $\dot{Q}(t) = -\lambda(\bar{p} - w)Q(t)$, $Q(0) = q$ and T is given by $Q(T) = r$, we easily obtain

$$\frac{\partial}{\partial r} V(q, \phi_r''') = e^{-\rho T} \left(\frac{w - c}{\rho} - \frac{1}{\lambda} \right) \leq 0 .$$

Thus, $V(\phi_\sigma''') \geq V^2$ with strict inequality unless $\sigma = \frac{1}{2}$.

14

We now recall the policies ϕ'_b defined above. Since $\sigma \leq \frac{1}{2}$ we have $H^1(\sigma) \geq H^2(\sigma)$ hence $V(\phi'_\sigma) \geq V(\phi'''_\sigma)$. The analysis in the proof of Lemma 3.2 shows that $V(\phi'_a) \geq V\langle\phi'_\sigma\rangle$ with strict inequality unless $a = \sigma$. Since $V^1 = V(\phi'_a)$ the proof is complete. $\quad\square$

Lemma 3.8. *If $\sigma, a \geq \frac{1}{2}$ then $V^2 \geq V^1$ with strict inequality unless $a = \frac{1}{2}$.*

Proof: Consider the policies

$$
\phi''''_r(q) := \begin{cases} m_0 = 1 \quad m_{\bar{p}} = 0, & q < a \\[2mm] m_0 = 0 \quad m_{\bar{p}} = 0 & a \leq q \leq r \\[2mm] m_0 = 0 \quad m_{\bar{p}} = 1 & r < q \,. \end{cases}
$$

Since $a \geq \frac{1}{2}$ we have $H^2(a) \geq H^1(a)$, hence $V(\phi''''_a) \geq V^1$ and the inequality is strict unless $a = \frac{1}{2}$. As in Lemma 3.7, for $q > r$ we obtain

$$
\frac{\partial}{\partial r} V(q, \phi''''_r) = e^{-\rho T} \left(\frac{w - c}{\rho} - \frac{1}{\lambda} \right) \geq 0 \,.
$$

Hence $V(\phi''_a) = V(\phi''''_1) \geq V(\phi''''_a)$. By Lemma 3.3 we have $V^2 \geq V(\phi''_a)$. This completes the proof. $\quad\square$

Lemma 3.9. *Assume $a < \sigma$. Then there exists at most one point $x \in [a, 1]$ such that $V^1(x) = V^2(x)$. If $q \in [a, 1]$ satisfies $q < x$ then $V^1(q) > V^2(q)$ and if $q > x$ then $V^2(q) > V^1(q)$. For $q \in [0, a)$ we have $\mathrm{sgn}(V^1(q) - V^2(q)) = \mathrm{sgn}(V^1(a) - V^2(a))$.*

Proof: We will prove the last statement first. For $q \in [0, a]$ we have $V^1(q) = e^{-\rho T(q)} V^1(a)$ and $V^2(q) = e^{-\rho T(q)} V^2(a)$ where $T(q)$ is given by $e^{-\lambda w T(q)}(1 - q) = 1 - a$. Thus $\frac{V^1(q)}{V^2(q)} = \frac{V^1(a)}{V^2(a)}$.

If $\sigma \leq \frac{1}{2}$ then $V^1 > V^2$ by Lemma 3.7. We assume $\sigma > \frac{1}{2}$. Let $q \in [\sigma, 1]$, then we have $V^2(q) = \frac{(w-c)q}{\rho}$ and $V^1_q(q) = \frac{-\rho V^1(q) + (\bar{p} - c)q}{\lambda(\bar{p} - w)q}$. It follows that if $V^2(q) = V^1(q)$, then $\frac{w-c}{\rho} = V^2_q(q) > V^1_q(q) = \frac{1}{\lambda}$. We can now conclude that if $V^1(x) = V^2(x)$ for any $x \in [\sigma, 1]$ then $V^1(q) > V^2(q)$ for $q \in [\sigma, x)$ and $V^1(q) < V^2(q)$ for $q \in (x, 1]$.

Now, let $q \in [a, \sigma]$ and assume $V^2(q) \geq V^1(q)$. Then, $V^2_q(q) = \frac{\rho V^2(q) + cq}{f(0,q)} \geq \frac{\rho V^1(q) + cq}{f(0,q)}$. Since $V^1_q(q) = \frac{\rho V^1(q) - \bar{p}q + cq}{f(\bar{p},q)}$ we have $V^2_q(q) \geq V^1_q(q)$ if $\frac{\rho V^1(q) + cq}{f(0,q)} \geq \frac{\rho V^1(q) - \bar{p}q + cq}{f(\bar{p},q)}$, which reduces to $V^1(q) \geq H^1(q)$. Furthermore, equality can hold only if $V^2(q) = V^1(q) = H^1(q)$, which implies $q = a$. The Lemma now follows from Lemma 3.2. $\quad\square$

We are now ready to prove the main result.

Proof of Theorem 2.4. Let $V = \max(V^1, V^2)$. In the case $V = V^1$ or $V = V^2$ we see that V is sufficiently regular[2] to proceed along classical lines. Both V^1 and V^2 are easily verified

---

[2] V^1 and V^2 are in fact both C^2 but we require only C^1.

to be C^1. Lemmas 3.2 to 3.7 show that if $V = V^1$ or $V = V^2$ then V satisfies HJB everywhere on $[0,1]$. Classical "verification theorems" [FS] can be applied proving that $V = V^{opt}$. Since the policy ϕ^{opt} achieves V it is optimal. The classical verification theorems can be extended to prove optimality of V even when $V \neq V^1$ and $V \neq V^2$. However, we can also quote the viscosity theory: since V is piecewise C^1 (by Lemma 3.9), V is the viscosity solution to HJB and it is therefore equal to V^{opt}. $\square$

4. Analysis and Simulation of the γ-AR Model

In this section we collect some ideas concerning the analysis of the γ-AR model. Since we do not have a solution we will be somewhat informal. By Blackwell's Theorem we can restrict to stationary policies $P = \phi(Q, \tilde{P})$.

Assuming $V_\gamma^{opt} = V_\gamma^{opt}(Q_0, \tilde{P}_0)$ is C_1 at $(Q_0, \tilde{P}_0)$, it satisfies the Hamilton-Jacobi-Bellman equation there,

$$\rho V(q, \tilde{p}) = \sup_{p \in [0, \bar{p}]} \left\{ V_1 f(\tilde{p}, q) + V_2 \gamma(p - \tilde{p}) + q(p - c) \right\},$$

where V_i denotes partial differentiation with respect to the ith argument. Since the expression inside the brackets is linear in p we see that we can assume $p = \phi(Q_0, \tilde{P}_0) \in \{0, \bar{p}\}$ without loss of generality. However, in so doing we may be forced to generalize the notion of the solution to $\dot{\tilde{P}} = \gamma(P - \tilde{P})$ to admit weak solutions: if $(Q(t), \tilde{P}(t))$ tracks a boundary between $\{\phi = \bar{p}\}$ and $\{\phi = 0\}$ then $P(t)$ 'oscillates' between 0 and $\bar{p}$. In this case it is possible to redefine ϕ along the boundary so that a classical solution admits. In fact, if for some optimal $P(t)$ we have $P(t) \in (0, \bar{p})$, $\dot{Q}(t) \neq 0$ for $t \in (a, b)$ then we can apply the calculus of variations to

$$V(Q(t), \dot{Q}(t), \ddot{Q}(t)) = \int_0^\infty e^{-\rho t} (P(Q, \dot{Q}, \ddot{Q}) - c) Q(t) dt$$

to obtain a differential equation for Q on (a, b). (This is possible since we can express both $\tilde{P}$ and P as functions of Q, $\dot{Q}$, and $\ddot{Q}$.) This will in turn determine $P(t)$ on (a, b). Furthermore, such a 'tracking boundary' between $\{\phi = 0\}$ and $\{\phi = \bar{p}\}$ determines a solution to this differential equation. Simulations indicate that the optimal value function is piecewise C^1 in general. Corresponding to the point x in the ∞-AR solution there arises a 1-dimensional curve along which the gradient of V_γ^{opt} appears discontinuous (see Fig. 4.8.)

Our simulations results were computed as follows. To each discrete point z in the $(Q, \tilde{P})$ space we assign a value $v(z)$ and an price, $p(z)$. We let Q and $\tilde{P}$ evolve according to their differential equations with initial condition z and $P = p(z)$ over a short interval $t \in [0, \epsilon]$ so

that $(Q(\epsilon), \tilde{P}(\epsilon))$ is still close to z. We then evaluate,

$$\int_0^\epsilon e^{\rho t}(p(z) - c)Q(t) \; dt \; + \; v(Q(\epsilon), \tilde{P}(\epsilon))$$

where $v(Q(\epsilon), \tilde{P}(\epsilon))$ is computed by locally interpolating v. For each point z we compute the above for all possible choices of $p(z)$ and then assign to $v(z)$ the maximum attained among the choices for $p(z)$. Initializing $v = 0$ and iterating the above described computation over all z we obtain an increasing sequence of v which converge to an approximation to V_γ^{opt}.

In view of the results above we restrict $p(z) \in \{0, \bar{p}\}$ except on the set $\tilde{P} = w$ where we admit $p(z) = w$. The latter is admitted since we expect stationary points along this set, in view of our convergence results. (We experimented with allowing $p(z) \in \{0, w, \bar{p}\}$ in general, and the solutions presented here were only slightly modified with $p(z) = w$ in a neighborhood of the tracking boundary, as we would expect.)

The following table lists various sets of parameter values which we have used for simulation.

Parameter Set						a	b	c
	λ	ρ	c	w	$\bar{p}$		γ	
1	0.5	0.1	0.0	0.6	2.0	1.0	2.0	5.0
2	0.5	0.1	0.0	0.5	2.0	2.0	5.0	$\star$

Table 1: Parameter Values for Simulation

Fig. 4.1 shows both V^1 and V^2 for the ∞-AR problem parameter sets 1 and 2. Note that for parameter set 1 $V^2 > V^1$ while for parameter set 2 the two functions cross.

We have plotted optimal vector fields for parameter sets 1a,1b,1c,2a, and 2b in Figs. 4.2-4.6. The optimal price is overlaid at each point in the discrete space. Oscillatory limit cycles can be seen in Figs. 4.5 and 4.6 corresponding to the oscillating solution of the ∞-AR problem arising in parameter set 2. Surprisingly such a limit cycle also appears in Fig. 4.2, corresponding to parameter set 1a. This indicates that small values of γ favor oscillating limit behavior. In Figs. 4.3 and 4.4, corresponding to parameter sets 1b and 1c, the oscillating limit is absent. This is to be expected from our approximation results.

The limit cycle in Fig. 4.5 (parameter set 2a) has larger amplitude in q then that of 4.6 (parameter set 2b) as we would expect from our approximation results. In the limit $\gamma \to \infty$ this amplitude goes to 0.

All of the vector fields have some stable stationary set along the line $\tilde{P} = w$. This is also to be expected from our approximation results.

17

Figs. 4.7 and 4.8 plot the value function for two extreme values of γ for parameter sets 1 and 2. Here we see the smoothness of the value function and also the convergence to V_∞^{opt}.

Acknowledgements. The authors thank Hsueh-Ling Huynh and Peter Linhart for many viscous discussions.

18

References

[FS]   Fleming, W. H. and Soner, H. M., *Controlled Markov Processes and Viscosity Solutions*, Springer-Verlag (1993).

[Ga]   Gamkrelidze, R. V., *Principles of Optimal Control Theory*, Translated by U. Makowski, Plenum Press, New York (1978).

[Li]   Lions, P. L., "Optimal Control and Viscosity Solutions," in *Recent Mathematical Methods in Dynamic Programming*, Springer-Verlag (1985).

[LS]   Lions, P. L., and Sougandis, P. E., "Differential Games, Optimal Control and Directional Derivatives of Viscosity Solutions of Bellman's and Isaacs' Equations," SIAM J. Control and Optimization, Vol. 23, No. 4, pp. 566–583, July (1985).

[CHL]   Crandall, M. G., Ishii, H., and Lions, P. L., "User's guide to viscosity solutions of second order partial differential equations," Bulletin of the AMS, Vol. 27, No. 1, pp. 1–67, July (1992).

[RR]   Radner, R., "Viscous Demand," Stern School of Business, New York University, June (1996), (unpublished).

Parameter Set 1
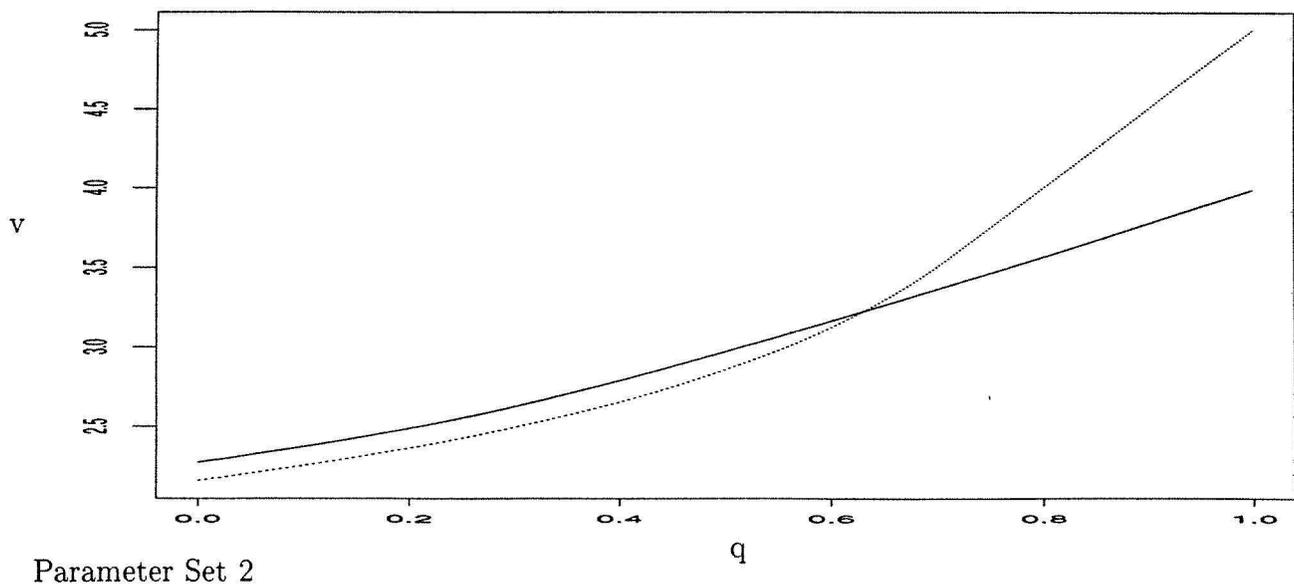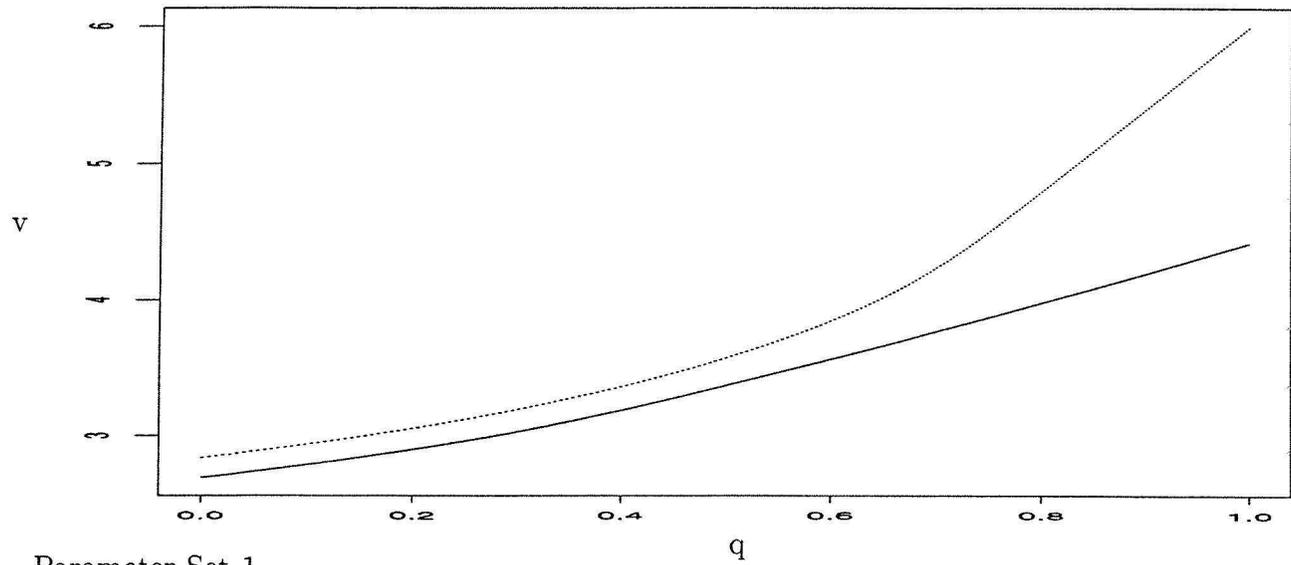
Parameter Set 2

Figure 4.1: Value Functions.
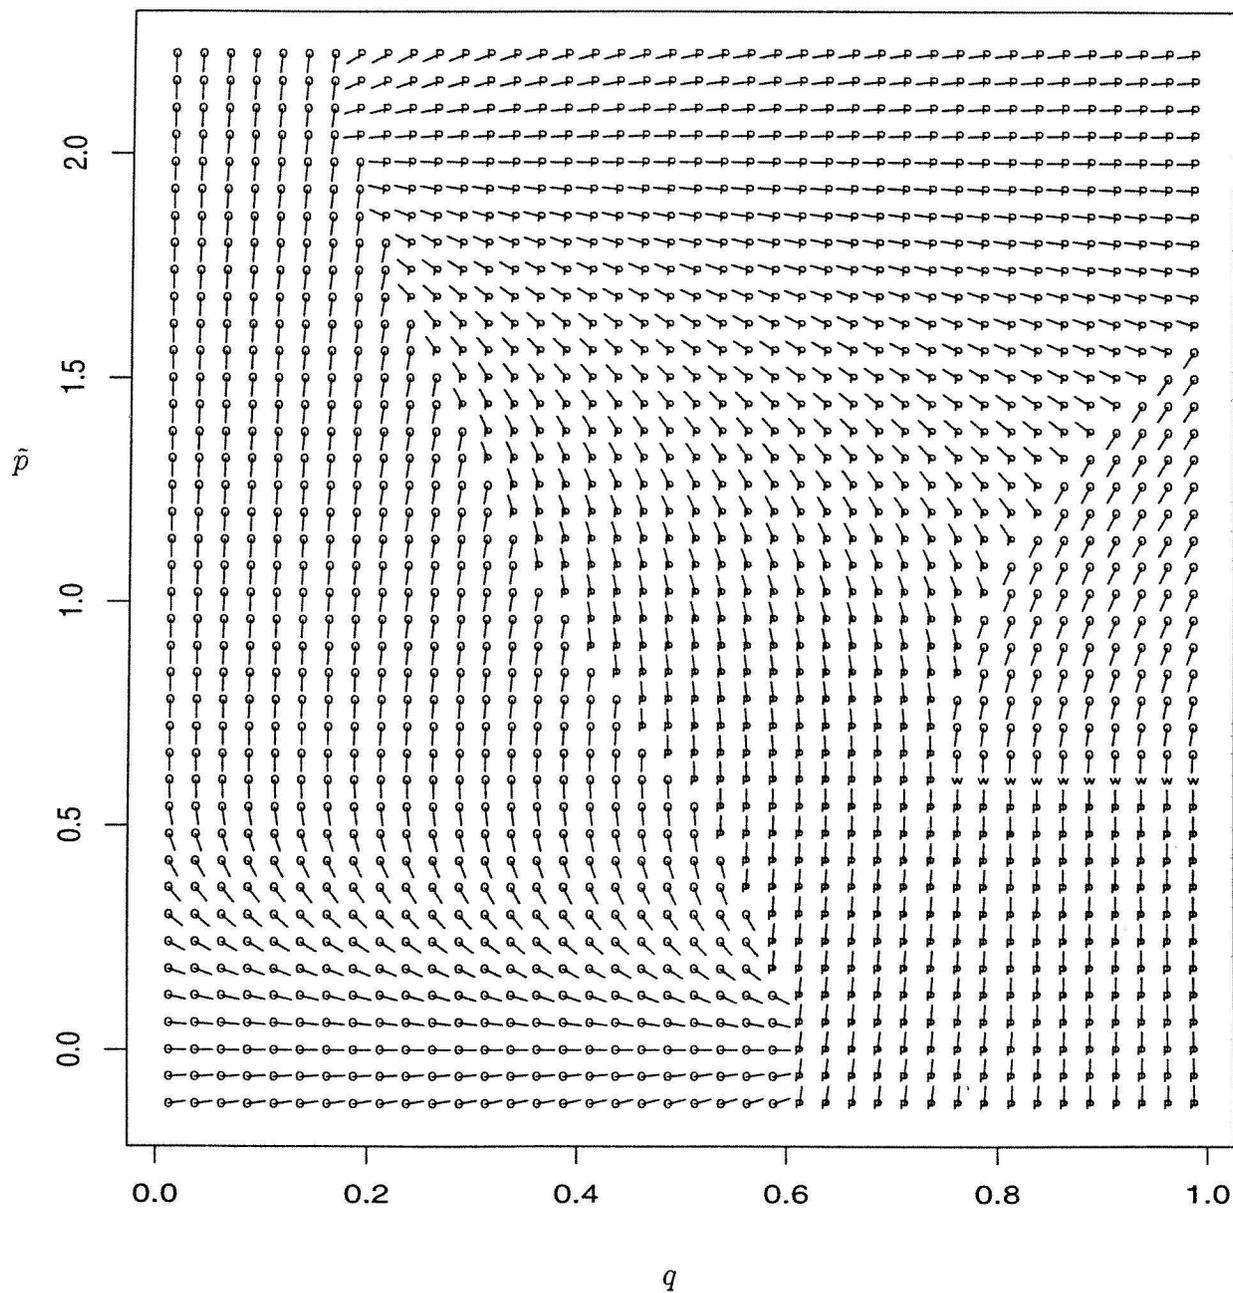
Figure 4.2: Optimal Vector Field for Parameter set 1a.
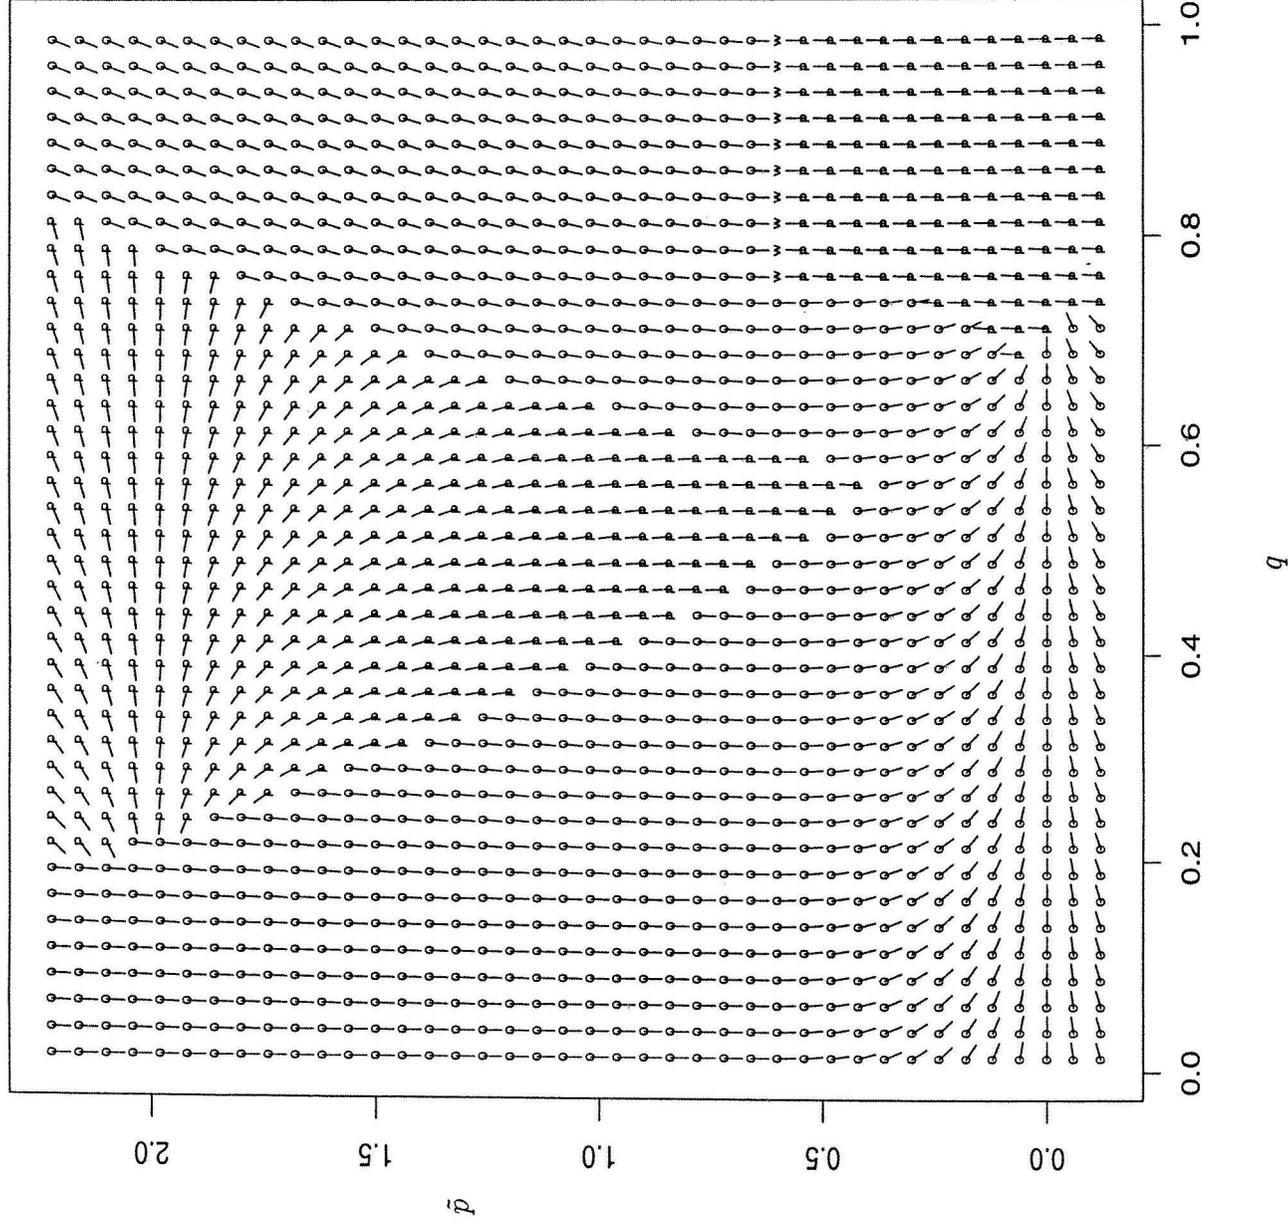
Figure 4.3: Optimal Vector Field for Parameter set 1b.
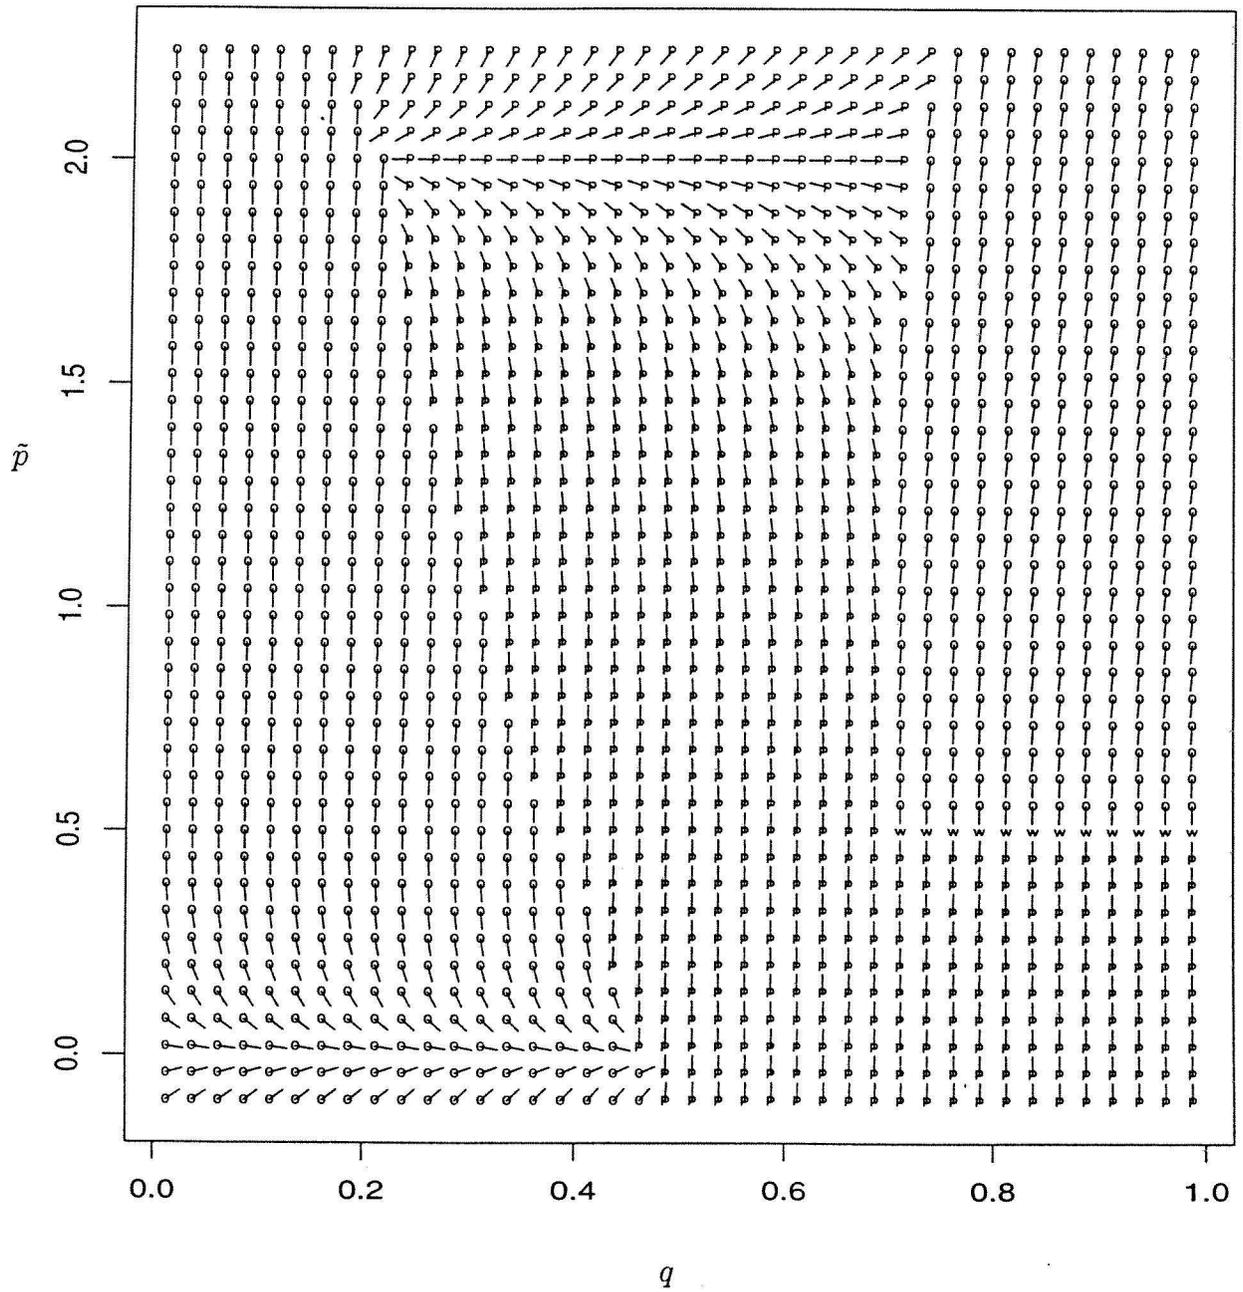
Figure 4.4: Optimal Vector Field for Parameter set 1c.
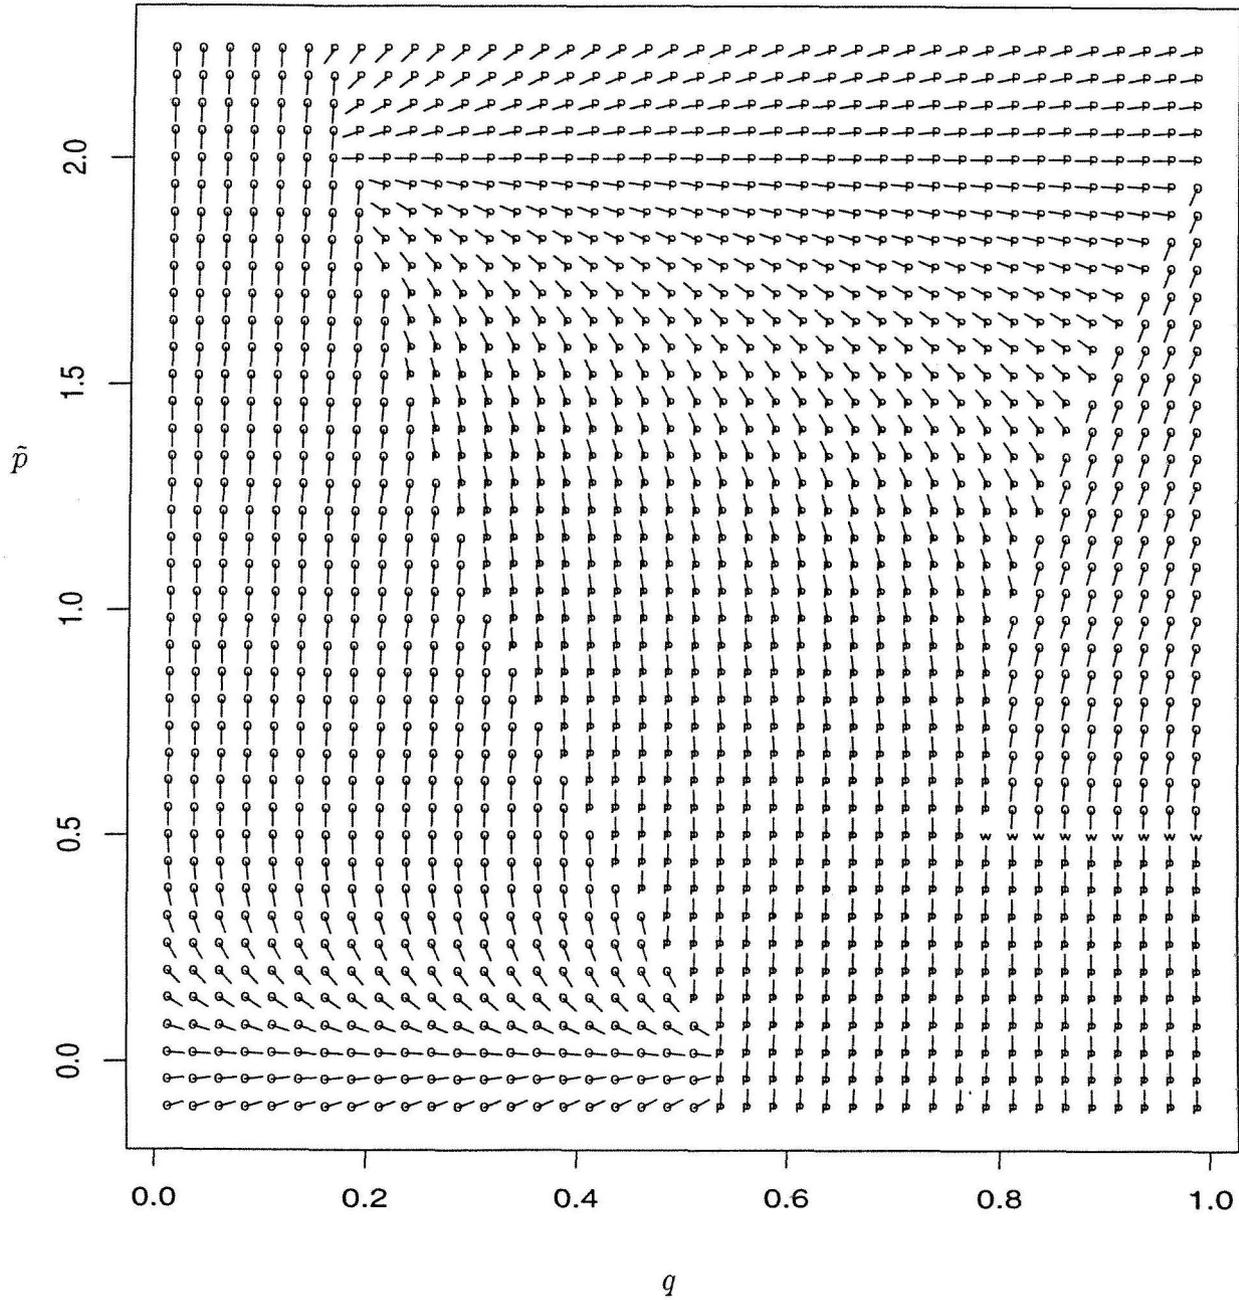
Figure 4.5: Optimal Vector Field for Parameter set 2a.
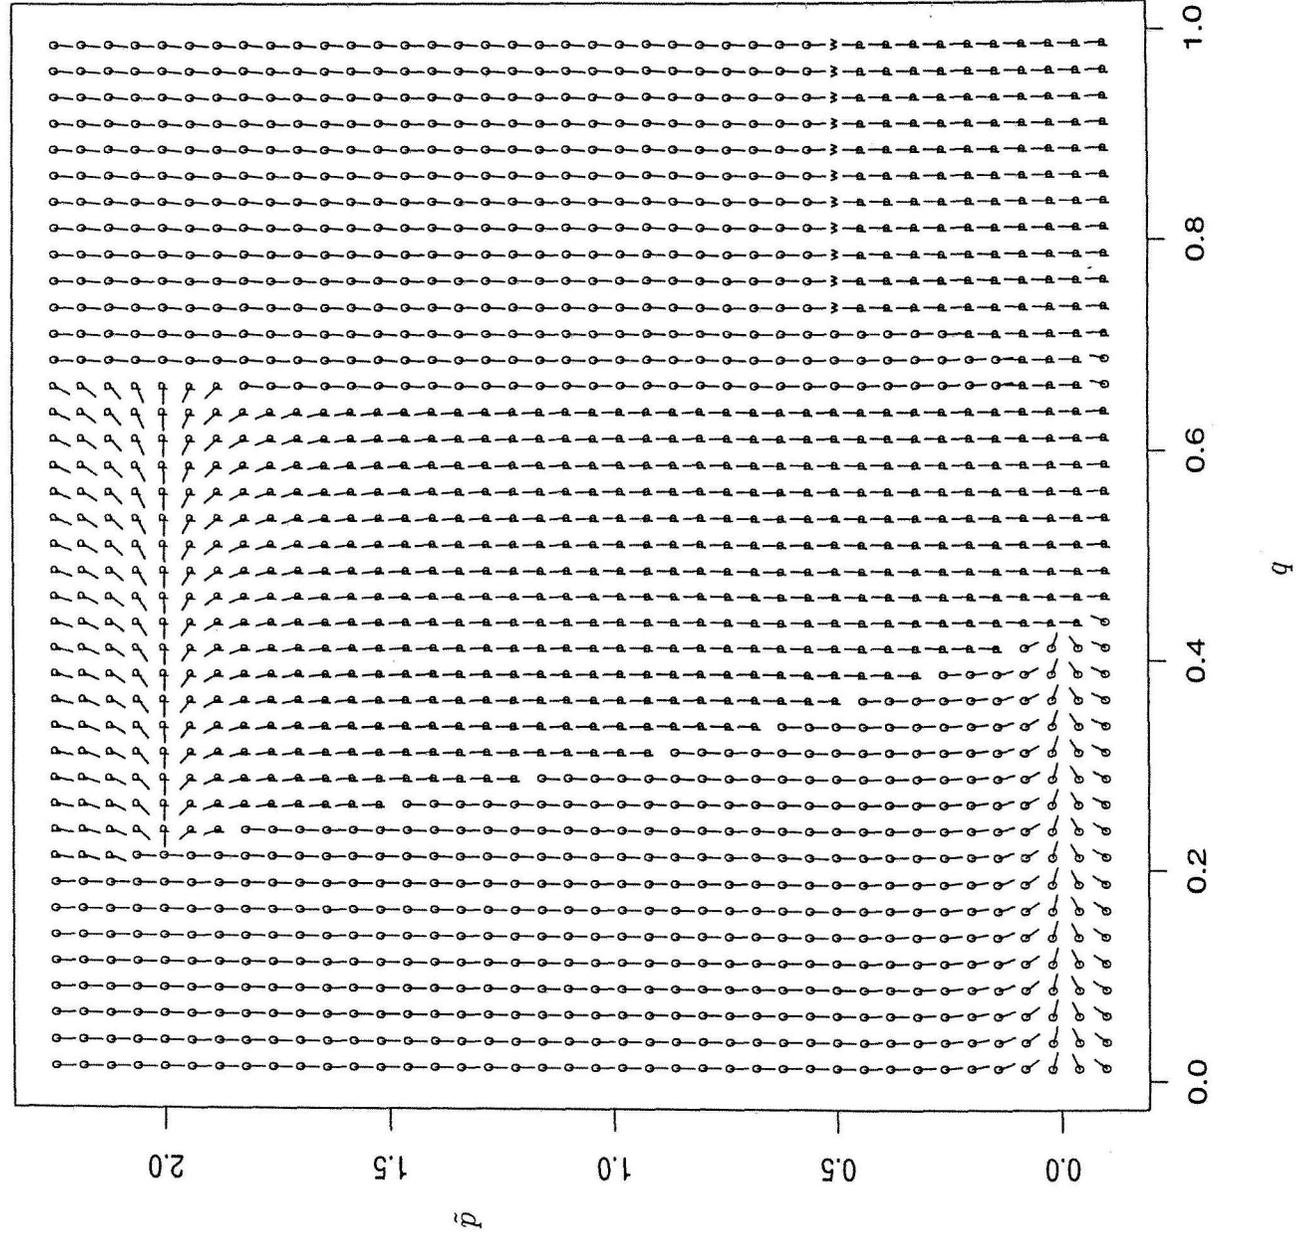
Figure 4.6: Optimal Vector Field for Parameter set 2b.
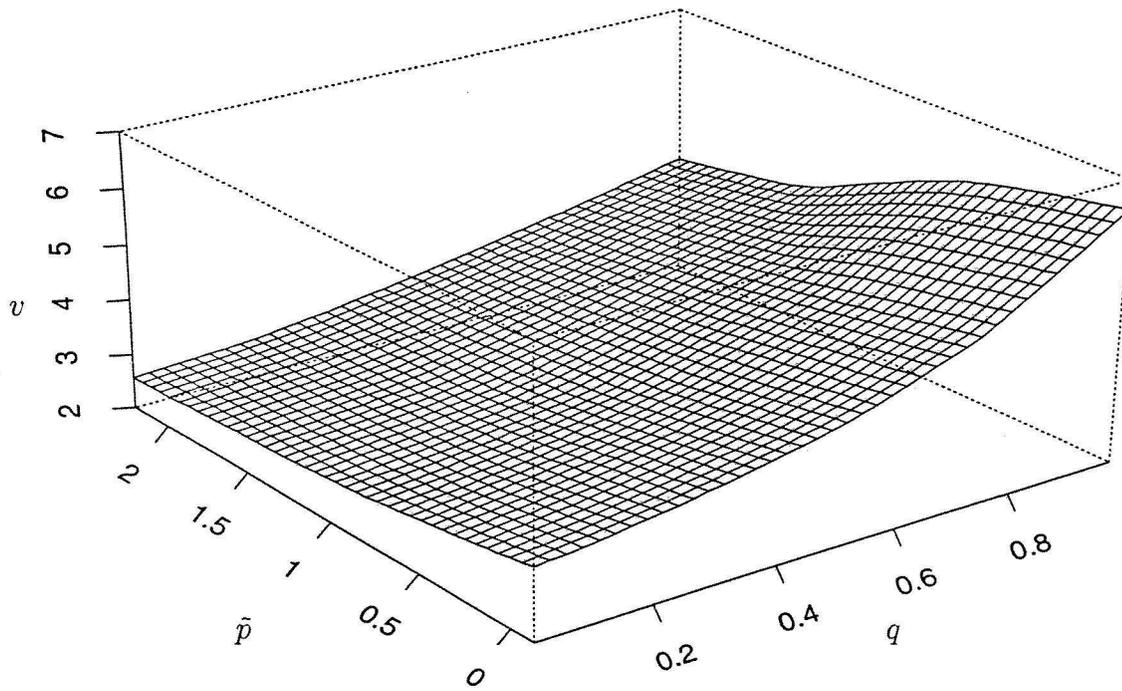
Parameter Set 1a
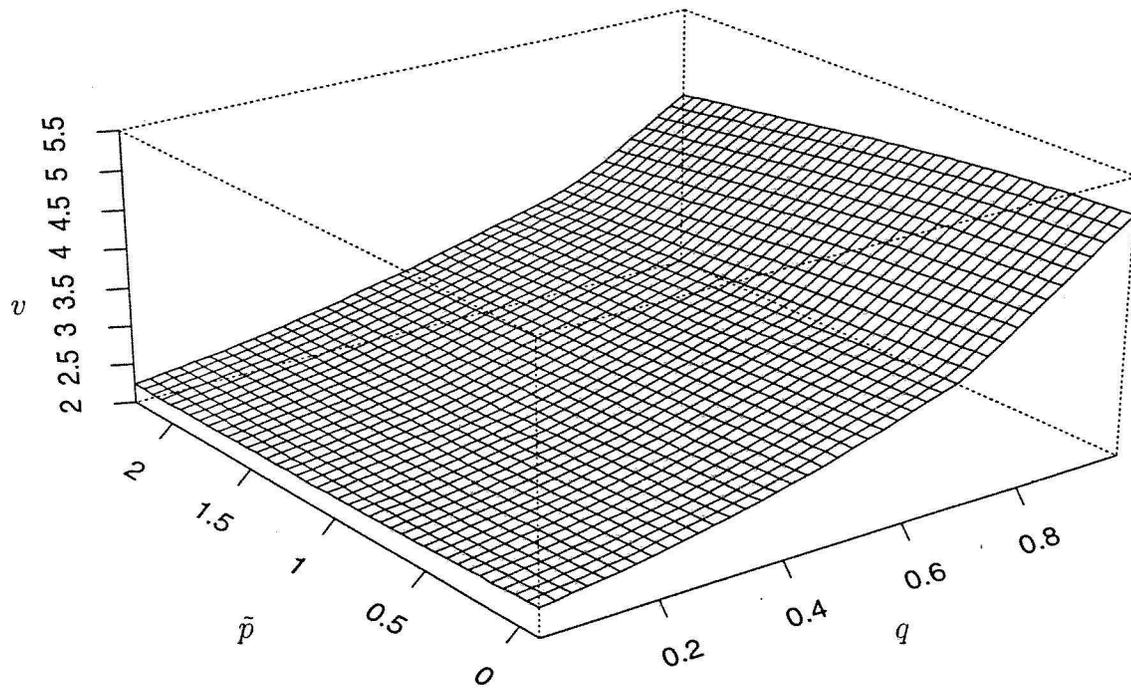
Parameter Set 1c

Figure 4.7: Optimal Value Functions: Dependence on γ.
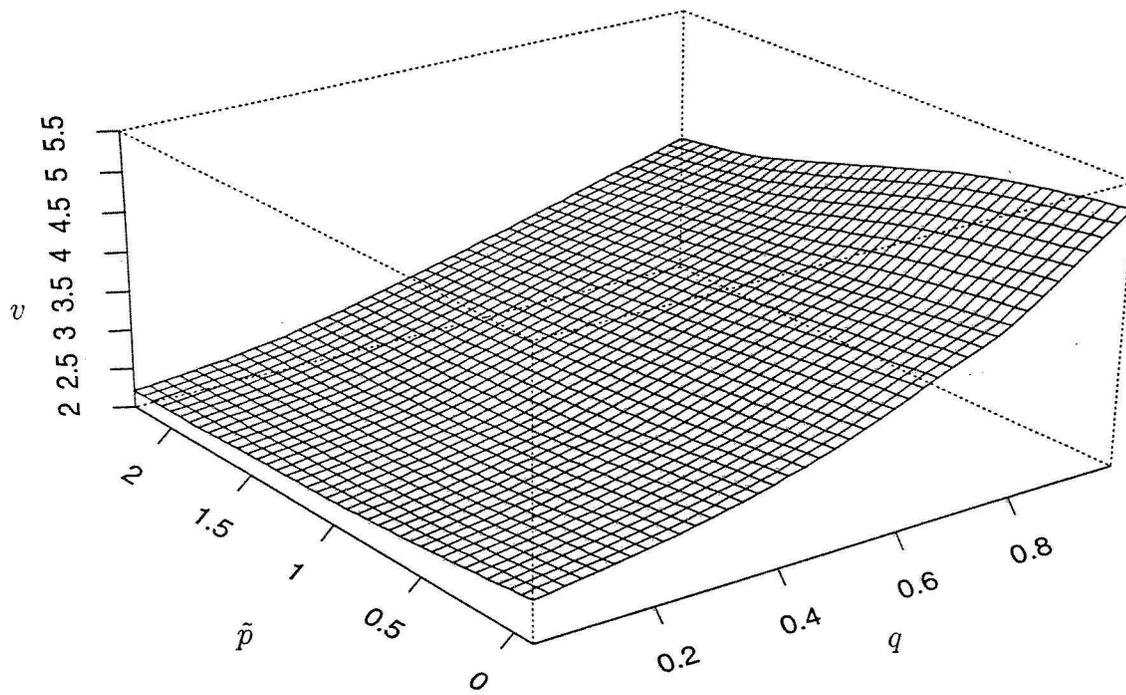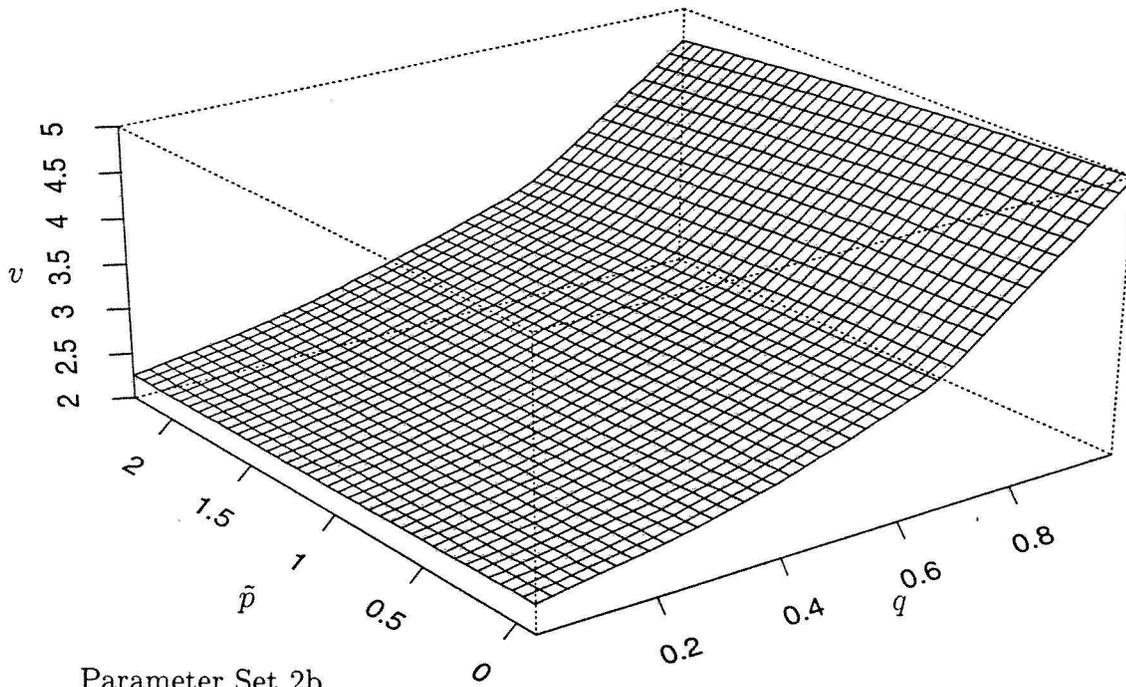
Parameter Set 2a

Parameter Set 2b

Figure 4.8: Optimal Value Functions: Dependence on γ.

27