QUANTIFYING AIRCRAFT AGILITY USING MINIMUM-TIME MANEUVERS

by

VERA ANN MARTINOVICH

Bachelor of Science, Aerospace Engineering
Iowa State University
(1988)

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Master of Science
in
Aeronautics and Astronautics
at the
Massachusetts Institute of Technology

September 1990

Signature of Author_____
Department of Astronautics and Astronautics
31 July 1990

Certified by_____
Professor Steven R. Hall
Thesis Supervisor, Department of Aeronautics and Astronautics

Accepted by_____
Professor Harold Y. Wachman
Chairman, Department Graduate Committee

QUANTIFYING AIRCRAFT AGILITY USING MINIMUM-TIME MANEUVERS

by

VERA ANN MARTINOVICH

Submitted to the Department of Aeronautics and Astronautics
on 31 July 1990 in partial fulfillment of the requirements
for the Degree of Master of Science in
Aeronautics and Astronautics

ABSTRACT

Aircraft agility is studied from an optimal control perspective. Maneuvers conventionally used to point an aircraft towards a target were compared to minimum-time maneuvers. Using the time-optimal maneuvers trimmed up to 48% from the pointing times. The difference comes from pitching and rolling simultaneously in a loaded roll. The nominal fighter model had limitations on roll rate, pitch rate, and engine spool time. Each of these limits was removed in turn and minimum-time pointing again conducted. Removing roll rate limits allowed the aircraft to point to targets below it about 10 deg farther in the same time. Similar advantages are gained by removing pitch rate limits. These improvements are independent of target position but are only for trajectories of moderate length. Spool time limits had virtually no effect on pointing times, although a small difference was noted for long trajectories. Finally, using an unloaded roll to compare agility was shown to be deceiving. This maneuver masks performance capabilities and makes the aircraft appear slower than it actually is.

Thesis Supervisor: Dr. Steven R. Hall

Title: Finmeccanica Assistant Professor of Aeronautics and Astronautics

Acknowledgments

I owe many debts for this thesis, and I would like to attempt to pay them here. I owe my gratitude to my friends who are also struggling through graduate school for understanding my frustrations in day-to-day life. Linda Abetz was a wonderful woman's voice in my mostly male world. Doug "Studs MacKenzie" MacMartin frequently saved my hide when the computers went on the fritz and was also with me on trips to Stefani's, the movies, and outlet mall shopping. Narendra "Bhat-man" Bhat helped me conquer Word 4.0 and always stopped by my desk to check up on me. My best friend Adam Sawicki was sometimes the only reason I went running when I felt like staying home and eating Ding-Dongs. He shared my problems and my joys and was the main reason I didn't lose my mind at the 'Tute.

There are several people I must thank for professional contributions to my thesis. Jim Purdon of the MIT Supercomputer Facility showed me how to optimize my program to run in 15 seconds instead of 15 hours. In fact, the whole Supercomputer staff were some of the most helpful people I have ever encountered.

Captain Al Lawless at Edwards AFB gave me the pilot's perspective on agility and answered all of my sometimes silly questions on realistic performance limits and maneuvers.

I will always be grateful to Dagfinn Gangsaas, my supervisor at Boeing Advanced Systems, for allowing me to be on leave of absence for two years to finish my degree. He was forever patient when I called to tell him there were more delays.

People always get schmaltzy at the end of the acknowledgments and thank their families for all of the support. I will be no exception. My husband, Scott Gremmert, and my dog, Belle, waited patiently for two years in Seattle for their wife and mom to come home. I am humbled every time I think of their sacrifice and realize there is no way to repay them. My father, mother, and sister also deserve my gratitude for their long-distance support. Thank God it's ending; I can't afford the phone bills any longer.

This research was supported by Contracts DLH315256, DLH404167, and DLH418482 from The Charles Stark Draper Laboratory, Inc. All of the simulations were performed on the MIT Supercomputer Facility CRAY-2.

Vera Martinovich
Cambridge, MA
08 July 1990

Contents

List of Figures

List of Tables

Nomenclature

SYMBOL	DEFINITION
A	matrix the columns of which are the linear constraint gradients in an optimization problem
A/B	afterburner thrust, lbs
B	vector containing functions of **x** and **u** to be bounded
b	vector containing functions of **z** to be bounded; vector of constant constraints placed on **z**
C_D	drag coefficient, dimensionless
C_L	coefficient of lift, dimensionless
$C_{L_{max}}$	maximum coefficient of lift, dimensionless
c	constraints placed on **z**. May be inequality, equality, linear, or nonlinear constraints
D	crossrange needed to execute a 180 deg heading change
DT	product of crossrange and time to make a 180 deg heading change
f	equations describing the dynamics of **x**
g	gradient (vector of first derivatives) of the cost function
g_{max}	maximum load factor, g's
H	Hessian matrix (matrix of second derivatives) of the cost function
h_e	specific energy, ft. Sum of potential and total energy divided by aircraft weight
J	cost or objective function to optimize
L	Lagrangian function formed by augmenting the integral cost with the inner product of the Lagrange multipliers and the constraints
L	constant lower bound on **B**
L_{max}	maximum lift, lbs
l	constant lower bound on **b**
N_x	axial load factor, g's

N_z	longitudinal load factor, g's
P	integral cost, function of **x**, **u**, and t
P_s	specific excess power, ft/s. Time rate of change of h_e
Pstab	stability axis load acceleration, g's
p	integral cost, function of **z** and t
S	wing planform area, ft^2
TA	total angle through which the aircraft nose sweeps as it points to reach a target, deg. Second rotation in a 1-2-3 Euler system
t	time, s
t_f	final time, s
t_k	time to achieve first firing opportunity or time to kill, s
$t_{k/r}$	time to kill and recover lost speed, s
t_{RC90}	time to roll and capture a 90 deg bank angle change, s
t_o	initial time, s
U	constant upper bound on **B**
u	control vector; constant upper bound on **b**
x	state variable vector
V	airspeed, kts or ft/s
V_c	corner or maneuver speed, kts or ft/s
W	aircraft weight, lbs
z	one component of **z**
z	vector of independent variables including the states and controls at each segment endpoint of a trajectory, times at each endpoint, and any special user-defined parameters
z*	value of **z** yielding the minimum cost
α	angle of attack, deg; linear search distance in a quadratic programming problem

α^*	linear search distance minimizing the quadraticized cost in a quadratic programming problem
α_{max}	maximum angle of attack, deg
Δ	vector containing the difference of the state derivatives specified in **f** and the state derivatives found from a cubic spline interpolation of the states and controls within trajectory segments
Φ	terminal cost as a function of **x** and t_f
ϕ	bank angle about aircraft longitudinal axis, deg. Third rotation in a conventional 3-2-1 Euler system. Also terminal cost as a function of **z** and t_f
γ	flight path angle, deg
λ	Lagrange multiplier vector
μ	vector which converges to the Lagrange multipliers as the cost minimum is approached
θ	elevation angle between velocity vector and aircraft longitudinal axis, deg. Second rotation in a conventional 3-2-1 Euler system
ρ	air density, slugs/ft^3
σ	maneuver plane or bank angle about the velocity vector, deg. First rotation in a 1-2-3 Euler system
ψ	azimuth angle from true north, deg. East is 90 deg. First rotation in a conventional 3-2-1 Euler system
$\dot{\psi}$	turn rate, deg/s

1 Introduction

1.1 Motivation

Agility is one of the most important new areas of flight dynamics research. The focus is shifting to agility because new technologies are changing combat strategies. For instance, all-aspect missiles are eliminating the traditional tail chase and thrust vectoring and vortex management allow a fighter to operate at angles of attack well above stall.

These and other technologies are greatly shortening battle times. In the Vietnam era, for example, fighters maneuvered for long periods to get the positional advantage necessary to fire rear-aspect infrared missiles. In contrast, a future supermaneuverable fighter may be able to pull to 70 deg angle of attack, change heading 180 deg, and shoot the enemy in about 5 s [1]. This shifts the design emphasis from steady aircraft performance to dynamic, unsteady, agile performance. Thus airframe designers and tacticians are forced to rethink old ways of quantifying aircraft performance.

The early quantifiers described a fighter's point and sustained performances. Maximum load factor, maximum sustained turn rate, and thrust-to-weight ratio are all common examples. Energy-maneuverability (E-M) evolved next. E-M plots show specific excess power as a function of turn rate and can be easily used to compare aircraft. [2] gives a good description of these and other performance metrics.

None of these metrics is much good at predicting the transient maneuverability characteristic of modern and future combat aircraft. In some cases, they may be entirely misleading [2]. In short, these metrics do not quantify agility.

To quantify agility, one must first define it; and this has been done in many ways. [3] gives a good summary of four definitions. What they all boil down to is the "ability to carry out more state change activity per unit time" [4]. The Air Force Flight Test Center qualifies this definition by the phrase "with precision and control" [3]. Thus what is needed are new parameters describing the ability to change state rapidly and precisely.

Current agility research seems to be focused on heuristically proposing new agility metrics and then validating them by simulation. The following are three examples of such research:

1. Riley and Drajeske [5] have shown that Eidetics' torsional agility metric predicts the exchange ratio, that is, the number of blue kills for every red kill. Torsional agility is

defined as turn rate divided by the time to roll 90 deg and capture a given bank angle while turning at the given rate.

2. McAtee's dynamic speed turn plots [6] display airspeed bleed rate at each turn rate and acceleration capability at each airspeed. This information is more useful than that contained in "doghouse plots" showing turn rate vs. airspeed and makes it easier to compare aircraft. Dynamic speed turn plots show how much energy is left at the end of a maneuver to accelerate away or to trade for a positional advantage if returning to combat.
3. Cannon [7] makes use of the DT parameter, where DT is the product of crossrange and time needed to complete a 180 deg heading change. He uses trajectory optimization to show that aircraft on a minimum-DT, 180-deg-heading-change trajectory acquire the first shooting opportunity over aircraft on a minimum-time trajectory.

1.2 Research Summary

The drawback of propose-and-validate approaches to agility is that they presuppose a maneuver and do not allow one to distill new, possibly superior, metrics. Therefore, in this thesis, optimal control is used to find trajectories and no *a priori* agility metrics are used. A pilot's goal is assumed to be to point the fuselage at a target and remain there long enough to shoot, all in minimum time. This minimum time is taken as the measure of agility.

The point-and-shoot maneuver was chosen for its importance even without the new technologies discussed earlier. For instance, although we now have beyond-visual-range missiles, combat often degenerates into close-in dogfighting. In such tight battles, guns remain one of a pilot's most important weapons [8]. So close-in gun combat certainly makes good point-and-shoot capability valuable [6].

This research seeks the optimal trajectories to point the nose in minimum time to targets in the upper, lower, forward, and rear hemispheres. This was done by imagining the airplane to be in the center of a sphere and using a numerical trajectory optimization code to point its nose to a series of 108 test targets all over the sphere. The aerodynamic model used was that of a high-performance fighter similar to the F-15 with three performance limitations: roll rate, angle of attack rate, and engine spool time. The model and limitations are discussed more fully in Chapter 3.

The optimal maneuvers found were compared to the standard maneuvers taught to pilots as the best way to acquire the same 108 targets. The comparisons were done in two ways. First, time histories of important state variables such as speed and angle of attack

were kept for each run. Second, the optimal times and the standard times to complete each target acquisition were displayed together on a polar contour plot as functions of maneuver plane and total angle through which the nose sweeps. The overlaid curves allow easy determination of the time advantages of optimal flight.

Agility can be resolved into components of pitch, roll, and longitudinal agility [3]. To investigate fighter agility further, the baseline model with performance limitations was modified three times. One of the limitations was removed in each new version, making one agility component infinite in each new model. Again the models were tested in the 108 point-and-shoot trials. Time histories were kept and contour plots constructed. Of course it is unrealistic to expect an aircraft to have instantaneous engine spool time or undiminished roll and pitch capability as modeled, but the results for the three new cases bound the performance enhancements possible from improvements in each agility component.

Finally it is shown that using *a priori* maneuvers to test agility may hide real agility advantages. The model with unlimited pitch rate capability was tested again, this time using standard maneuvers. The results showed that the standard maneuvers concealed the true pitch capability of the model.

1.3 Contributions of the Thesis

This thesis provides four main contributions to the field of agility research:

1. Changes to standard maneuvers to achieve faster pointing times are discussed.
2. The relative importance of the three agility components is shown.
3. A method of displaying time-to-point data for easy aircraft comparison is provided.
4. Optimal control is promoted as a way of studying agility.

1.4 Outline of the Thesis

Chapter 1 contains the introduction, which gives the motivation for this research, outlines the research plan, and lists the contributions the thesis makes to the study of agility.

Chapter 2 discusses the trajectory optimization code by outlining the algorithm used and the program features. It also gives a brief history of agility metrics.

Chapter 3 discusses the research plan in greater depth, describing the aerodynamic model used and explaining the time contour plots more fully.

Results, including time histories, polar plots, and the actual time data, are in Chapter 4.

Chapter 5 ends the thesis with conclusions and recommendations.

2 Theory

This chapter discusses the theory behind the numerical optimization technique used to find the minimum-time trajectories. It also gives a history of agility measurement from the World Wars to the present and describes proposed metrics for future fighters.

2.1 Numerical Optimization

The trajectory optimizations were performed using a program called Optimal Trajectories by Implicit Simulation (OTIS) developed by the Air Force Flight Dynamics Laboratory. An overview of the optimization procedure used is given below.

An optimization problem is one in which a cost function is to be minimized subject to certain constraints. Mathematically it is written as

$$\begin{aligned} \min \quad & J = \Phi(\mathbf{x}(t_f), t_f) + \int_{t_o}^{t_f} P(\mathbf{x}(t), \mathbf{u}(t), t)\, dt \\ \text{subject to:} \quad & \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) \\ & \mathbf{L} \le \mathbf{B}(\mathbf{x}(t), \mathbf{u}(t)) \le \mathbf{U} \end{aligned} \tag{2.1}$$

The same formulation can be used to maximize a function by minimizing its negative.

The OTIS user begins by choosing a cost function. The easiest cost computationally is the terminal value of a state, control, or the time. For example, in this research, the final time to complete a pointing maneuver is the cost function. Other costs might be final weight or reentry vehicle surface temperature. A typical function to be maximized is range.

Next the constraints must be chosen. Some of these constraints, such as α_{max}, will be discussed in Section 3.3, "Performance Limitations." More important constraints are the equations of motion. OTIS is a three-degree-of-freedom program, meaning only the three force equations that define the movement of the aircraft center of gravity are used [9,10]. The OTIS user must choose an appropriate reference frame in which to integrate these three equations. The possible choices are Cartesian, flight path, and body coordinates. Next initial conditions corresponding to the chosen frame are selected. Additional constraints not specifically listed by the user will be discussed later.

Attitude controls are selected next. These need not necessarily be in the same reference frame as the equations of motion. For instance, in this research flight path equations, written in terms of airspeed (V), heading (ψ), and flight path angle (γ), were used. The controls, angle of attack (α) and bank angle (ϕ), are in the body frame. Forces corresponding to these controls are found through user-input tables that give force coefficients as functions of the controls and the flight condition. It is these forces that must be transferred to the flight path frame, and this is easily done with transformation matrices [11]. Once the transformation is complete, the equations can be integrated.

Rather than integrating the equations explicitly as is customary, implicit integration is performed [12,13]. To use this method, one first divides the trajectory into any number of equal or unequal segments. Nine equal segments were used in this research. Next "defect equations" are written for each segment center. A defect is the difference between the state derivative found from Equation (2.1) and the state derivative found from a cubic spline interpolation of the states and controls within segments. When these defects are set equal to zero, they form the additional non-user-written constraints mentioned earlier in this section.

The optimization problem is slightly modified from the form in Equation (2.1). It is now:

$$\begin{aligned} \min \quad & J = \phi(\mathbf{z}(t_f),t_f) + \int_{t_o}^{t_f} p(\mathbf{z}(t),t)\, dt \\ \text{subject to:} \quad & \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t),\mathbf{u}(t),t) \\ & \mathbf{l} \le \mathbf{b}(\mathbf{z}(t)) \le \mathbf{u} \\ & \Delta = 0 \end{aligned} \tag{2.2}$$

where $\mathbf{z}$ is a vector of independent variables, including the states and controls at each segment endpoint, the time of each segment endpoint, and any special user-defined design parameters. The bounds and the integrand are now written in lower case to signify that they are functions of $\mathbf{z}$, not $\mathbf{x}$ and $\mathbf{u}$. To the user, the cost function appears as in Equation (2.1).

After posing the optimization problem as in Equation (2.2), the actual numerical solution can begin. This is done in an OTIS subprogram called NPSOL written by the Systems Optimization Laboratory at Stanford University [14,15]. This subprogram uses a sequential quadratic programming (SQP) algorithm to solve (2.2) [16]. The user first guesses an initial trajectory, which could be quite different from the optimal trajectory

depending on the specific problem. The corresponding nominal control history is found from the initial guess, and the nominal control and state yield the nominal cost.

The solution proceeds differently depending on the presence and character of the constraints. The case of no constraints will be discussed first. This situation is usually handled by reducing the solution procedure to a quadratic programming (QP) problem. The first QP step is to find a suitable descent direction, that is, a **z** vector which provides a lower cost than the nominal. The negative gradient direction is a possible choice. The QP method approximates the cost as a quadratic, an especially good assumption as the minimum is approached. This quadratic approximation allows the use of a quasi-Newton search for a descent direction. Using this search, an approximation of the cost Hessian is formed and is updated at every iteration. This means problems can occur if positive-definiteness of the Hessian cannot be maintained [17,18].

During the second step of the QP method, the algorithm searches for a minimum along the descent direction found previously. This is an example of univariate minimization. Suppose the distance along this descent vector is called α. $J(\alpha)$ is approximated as a quadratic cost and the QP method tries to find the minimum cost J_1 located at α^*. J_0 is the current best guess of the optimal cost. This QP step is illustrated in Figure 2.1.

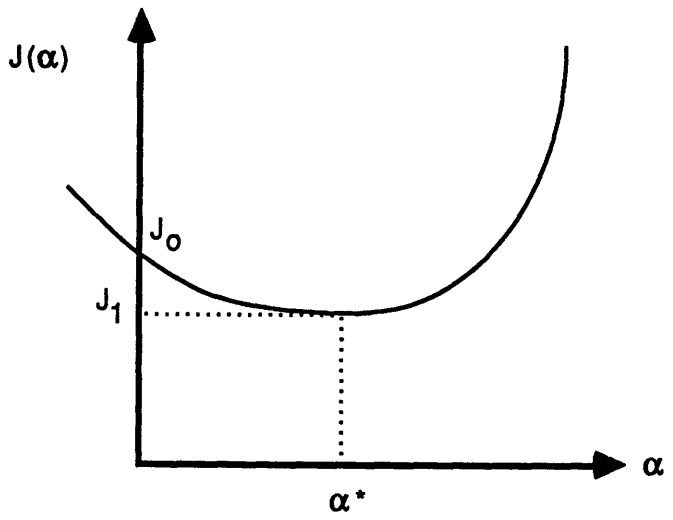

Figure 2.1. Linear search for the function minimum along the descent direction

Once J_1 is found, it becomes the current best guess of the optimal cost. Then a new descent direction for this new J_0 is found. The process terminates when a suitably small cost gradient is found.

Now the QP process can be illustrated for a fictitious problem in two dimensions with no constraints. This is shown in Figure 2.2 for a truly quadratic function. The curves shown are contours of constant cost. The independent variables are z_1 and z_2. J_1 and J^* are the results of two QP iterations to find the minimum along the two search directions shown.

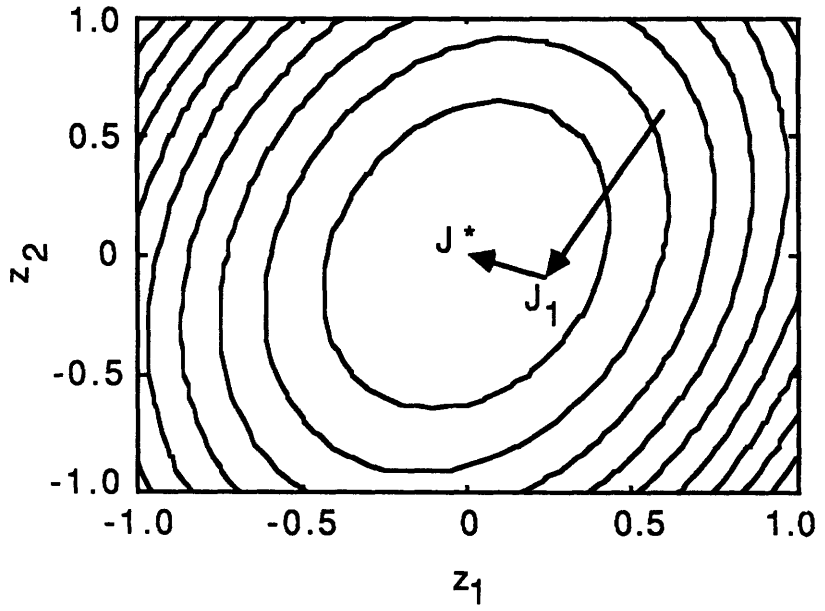

Figure 2.2. Quadratic programming solution in two steps for a truly quadratic function

The situation is considerably more complicated when constraints are present. There are two types of constraints, equality and inequality, with mixtures of both types allowed. These two classifications can be further divided into linear and nonlinear constraints. The worst case is nonlinear inequality constraints, which are dominant in flight dynamics. For example, limits on the roll rate as a function of angle of attack are usually nonlinear (see Section 3.3). Fortunately NPSOL and the SQP method are particularly suited for handling such constraints.

Inequality constraints are either active or passive: the minimum lies either on the constraint or off it on the appropriate side. Active constraints can therefore be treated as equality constraints. The problem is that the active constraints are not known *a priori*, and some sort of strategy to decide whether a particular constraint is active must be developed. For the immediately following paragraphs, it is assumed that an active set of constraints has been found; so all discussion holds for equality constraints. Later an active set strategy will be discussed.

We will first consider linear equality constraints of the form

$$\mathbf{c}(\mathbf{z}) = \mathbf{A}^T\mathbf{z} - \mathbf{b} = \mathbf{0} \tag{2.3}$$

with the columns of **A** linearly independent. Any feasible direction $\Delta\mathbf{z}$ from a stationary point must be tangent to the constraints, that is, it must lie in the tangent hyperplane formed by the intersection of all constraints (Figure 2.3(a)). Since the m columns of **A** are the gradients of each of the m constraints, the feasible direction must be orthogonal to the columns of **A**. Mathematically this is written as

$$\mathbf{A}^T\Delta\mathbf{z} = \mathbf{0} \tag{2.4}$$

Next we expand the cost J about the minimum point $\mathbf{z}^*$:

$$J(\mathbf{z}^*+\Delta\mathbf{z}) = J(\mathbf{z}^*) + \Delta\mathbf{z}^T\mathbf{g}(\mathbf{z}^*) + 1/2\Delta\mathbf{z}^T\mathbf{H}(\mathbf{z}^*)\Delta\mathbf{z} + \text{H.O.T.} \tag{2.5}$$

where $\mathbf{g}(\mathbf{z}^*)$ is the cost gradient evaluated at the minimum, $\mathbf{H}(\mathbf{z}^*)$ is its Hessian matrix, and the higher-order terms (H.O.T.) are neglected. As long as the Hessian is positive-definite, a higher cost in the chosen $\Delta\mathbf{z}$ direction is assured if

$$\Delta\mathbf{z}^T\mathbf{g}(\mathbf{z}^*) = 0 \tag{2.6}$$

Otherwise there are possible choices of $\Delta\mathbf{z}$ which let $J(\mathbf{z}^*) > J(\mathbf{z}^*+\Delta\mathbf{z})$ and $\mathbf{z}^*$ is no longer a minimum [17].

Comparing (2.4) and (2.6) we see that $\mathbf{g}(\mathbf{z}^*)$ is a linear combination of the columns of **A**. This can be expressed as

$$\mathbf{g}(\mathbf{z}^*) = \mathbf{A}\lambda$$

or

$$\mathbf{g}(\mathbf{z}^*) - \mathbf{A}\lambda = 0 \tag{2.7}$$

where λ is the vector of Lagrange multipliers.

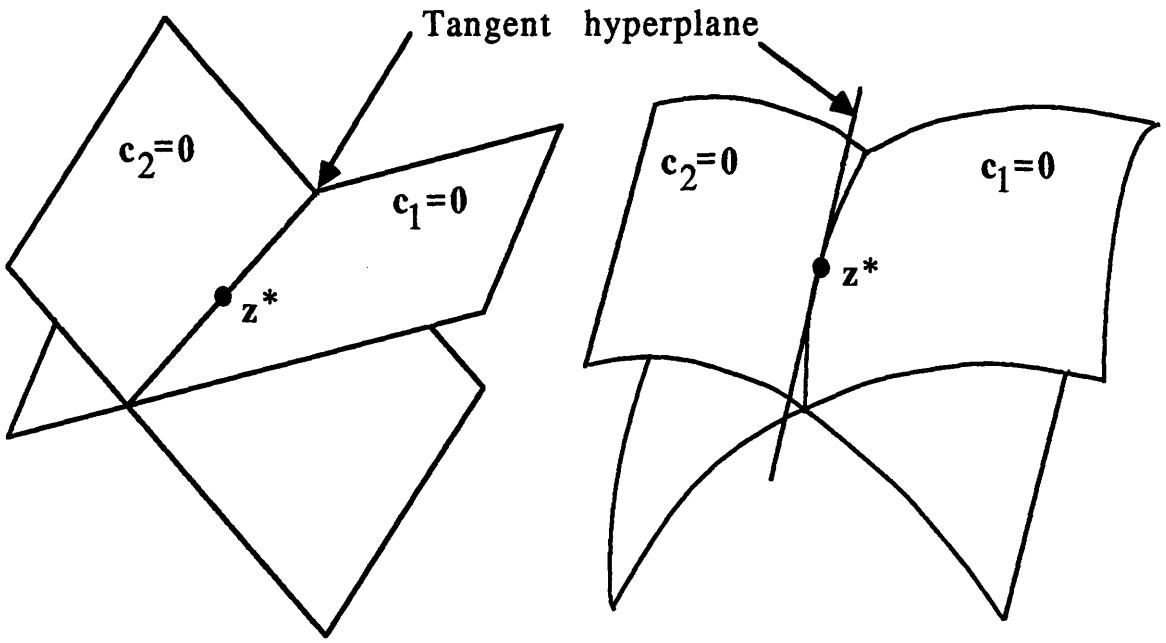

Figure 2.3. Tangent hyperplanes for (a) linear constraints and (b) nonlinear constraints

Now we have a system of simultaneous linear equations to be solved:

$$\mathbf{g(z)} - \mathbf{A}\mu = 0 \tag{2.7}$$

$$\mathbf{c(z)} = \mathbf{A}^T\mathbf{z} - \mathbf{b} = \mathbf{0} \tag{2.3}$$

The solution $(\mathbf{z},\mu)$ of this system is not guaranteed to be the local minimum $(\mathbf{z}^*,\lambda)$, only a stationary point.

To combat this difficulty, we try a different approach: minimize the Lagrangian function in the hyperspace in which the second-order sufficient conditions for a strong minimum hold. The Lagrangian function is defined as

$$L(\mathbf{z}) = p(\mathbf{z}) - \lambda^T\mathbf{c(z)} \tag{2.8}$$

where $p(\mathbf{z})$ is the integrand of (2.2) and λ is the Lagrange multiplier vector as before. The gradient of this Lagrangian function is the left-hand side of (2.7). This means the Lagrangian function has an unconstrained stationary point where the cost function has a constrained stationary point. Thus, finding a stationary point of (2.8) is the same as solving (2.7). The second-order sufficient condition is that the Hessian matrix of the

Lagrangian function must be positive-definite in all planes except that orthogonal to the tangent hyperplane formed by the intersection of the constraints at the minimum [16]. The orthogonal hyperplane is excluded because a feasible direction cannot lie in it.

The situation is not much more difficult for nonlinear equality constraints. Again the feasible directions must lie initially along the tangent hyperplane at the optimum as shown in Figure 2.3(b). Therefore the necessary conditions for a minimum are the same as in (2.7) and a Lagrangian minimization solution method can be used.

Projected Lagrangian methods, of which SQP is a subset, are designed to handle such nonlinear constraints in the manner just described. The idea here is to solve a series of minimization subproblems subject to linear constraints gradually approaching the tangent hyperplane that the nonlinear constraints form. After the initial guesses of $\mathbf{z}^*$ and λ have been made, the constraints are linearized about this guessed optimum and the Lagrangian function is minimized subject to these linear constraints. In this way, the subproblem solution is projected into the tangent hyperplane formed by the constraint approximations. The subproblem solution is the new guess of $\mathbf{z}^*$ and λ. The process repeats until the true solution is found. As the optimization progresses, the linear approximations become more accurate and eventually lie in the tangent hyperplane formed by the nonlinear constraints at the true minimum.

SQP simplifies the problem further by approximating the Lagrangian function as a quadratic, as well as linearizing the constraints, at every new guess of the optimal $\mathbf{z}^*$ and λ. Thus SQP is a series of QP problems described earlier. The SQP method is illustrated for a fictitious problem in Figure 2.4.

Until now we have assumed the absence of inequality constraints or the knowledge of which inequality constraints are active. However, most practical engineering problems have a mixture of equality and inequality constraints; and we do not know beforehand which inequality constraints affect the solution of the problem. There are several ways to deal with this difficulty, but the easiest is to choose a set of inequality constraints believed active and solve the subproblem using this set. The Lagrange multipliers of active constraints must be positive or zero [17] if the constraints are written as

$$\mathbf{c}(\mathbf{z}) \geq \mathbf{0} \tag{2.9}$$

This fact can be used to check the validity of the active set choice.

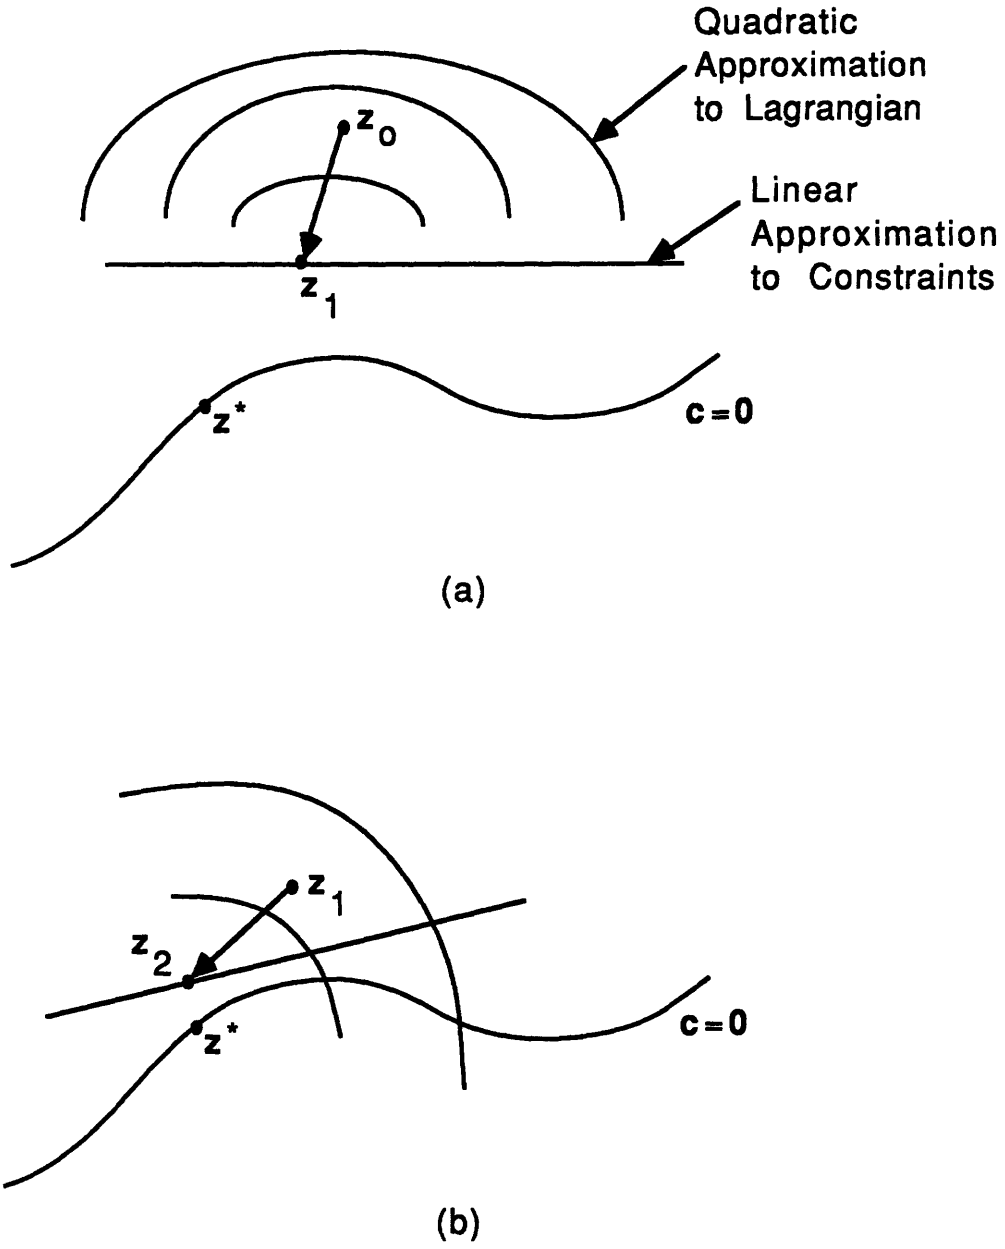

Figure 2.4. Demonstration of the SQP solution technique. (a) shows the first QP subproblem. The constraints are linearized and the function quadraticized about the initial guess z_0. The minimum of this subproblem is z_1. In (b) the new initial guess is z_1, about which approximations of the constraints and function are found. The minimum of this subproblem is z_2, which approaches the actual minimum at z^*.

2.2 Agility Background

This section provides a more in-depth view of the history behind agility metrics than that given in the introduction. There really is not a unified "theory of agility" yet because the topic has only recently received much emphasis. Every researcher seems to have his or her own pet theory. So instead of a theoretical discussion, this section briefly describes specific metrics used throughout the history of aviation, from the early sustained or point performance metrics to the more complicated ones of today.

The so-called point metrics have always been popular, in the past for their accuracy in predicting battle outcomes [2] and today for their simplicity. They are designated "point" because they describe aircraft static performance at a given flight condition. Combat in World Wars I and II was characterized by dogfights involving multiple slow-moving, fast-turning aircraft [19]. Two valid point metrics were therefore the maximum turn rate and minimum turn radius.

By the Korean era, aircraft cruised and climbed faster. The thrust-to-weight ratio (T/W) reflected an aircraft's ability to accelerate in a climb or on a straightaway.

The V-n diagram provided an early graphical comparison of aircraft. A typical V-n diagram is shown in Figure 2.5. Much important information comes from such a plot. The maximum speed and load factor are easily understood. The curved lines on the left represent the stall speed at each load factor. The corner speed is the speed at which the stall line intersects the structural load limit line. It is at this airspeed that the maximum turn rate and minimum turn radius are achieved [20].

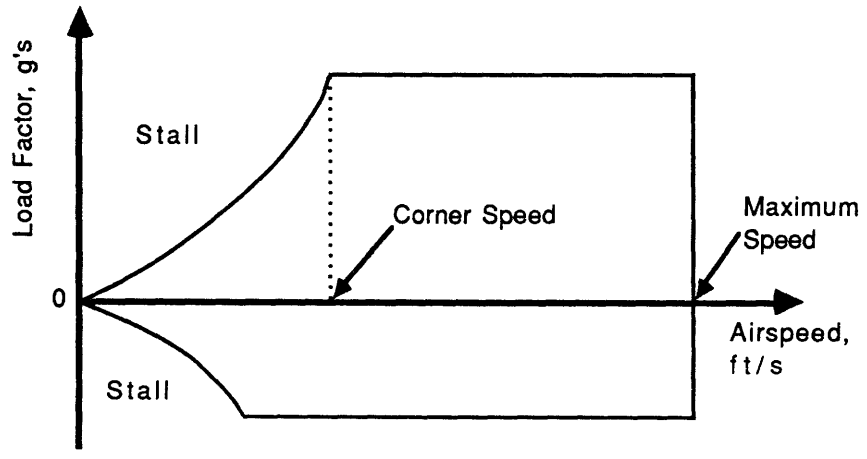

Figure 2.5. Typical V-n diagram [20]

One of the chief disadvantages of the V-n diagram is that it holds only for an aircraft in a constant altitude, constant airspeed coordinated turn [20]. With the emphasis shifting towards dynamic maneuvers, this graphical agility comparison is losing its validity.

To make up for some of the V-n diagram deficiencies, the energy-maneuverability (E-M) concept was introduced in the 1960s. E-M refers to the ability of an aircraft to change the potential and kinetic energies which comprise its energy state. Thus it accounts for changing speed and altitude. If the energy in ft-lbs is divided by the aircraft weight, a measurement called the specific energy results. The rate of change of this specific energy is called specific power, P_S. When $P_S > 0$ the aircraft can accelerate or climb. Steady, level flight occurs at $P_S = 0$. A negative P_S means deceleration or a loss in altitude [20].

One of the most common ways of displaying P_S information is in a plot such as Figure 2.6, in which the P_S of two competing aircraft are plotted against their turn rates at a given flight condition [2]. This gives a quick, uncluttered comparison of the relative strengths of each opponent.

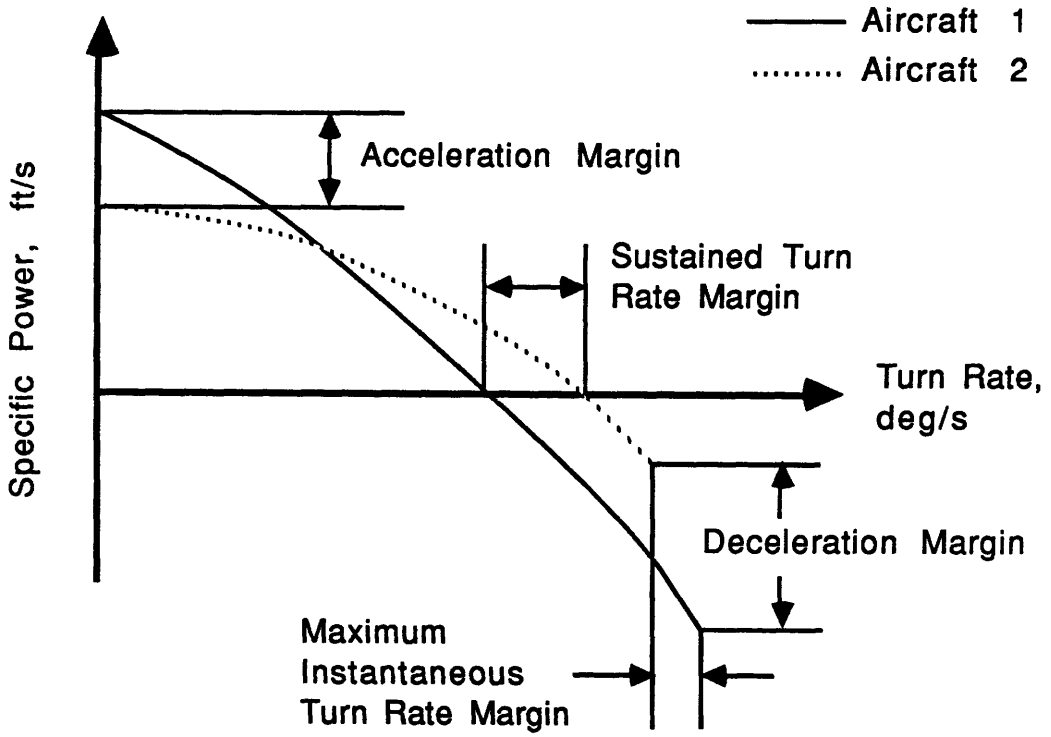

Figure 2.6. Typical energy-maneuverability comparison for two different aircraft for a given Mach number and altitude [2]

Another common E-M plot is the so-called "doghouse plot," shown in Figure 2.7. All of the V-n diagram information (maximum g, corner speed, and maximum speed) is available here, too. The constant P_S lines provide additional information about dynamic maneuvers. The peak in the $P_S = 0$ contour tells us the maximum sustained g. The other P_S contours tell us how fast the aircraft is losing or gaining energy for a specific condition. The doghouse does not tell a pilot how long he can maintain his P_S advantage as his flight condition changes.

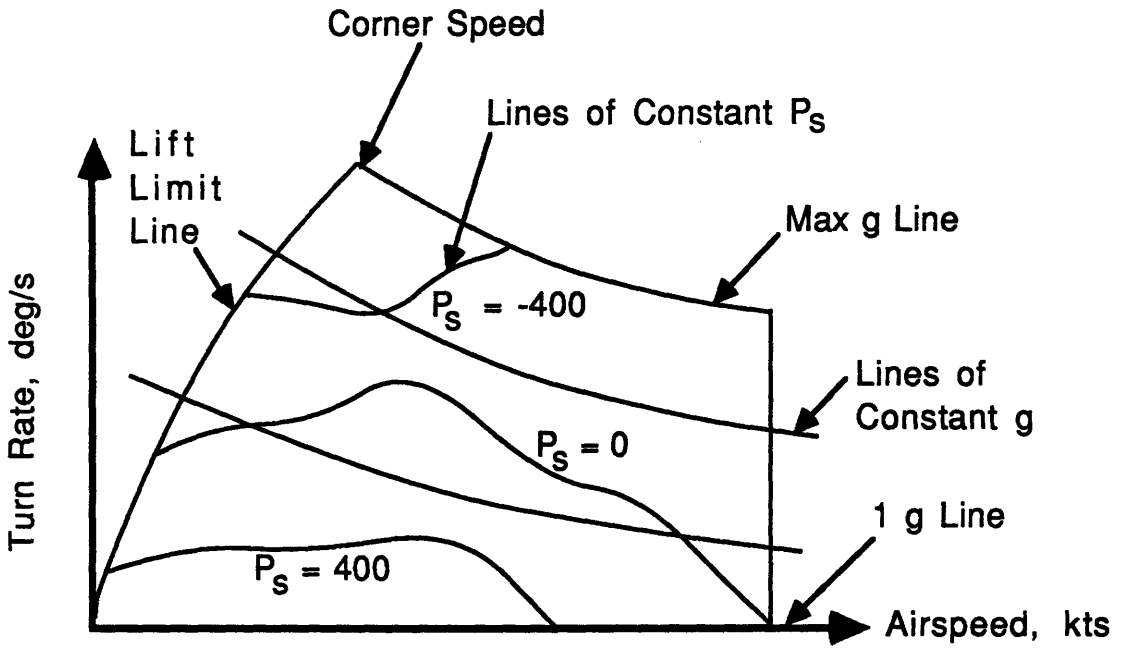

Figure 2.7. Typical doghouse plot [19]

The E-M plots predict engagement outcomes well when sustained maneuvers are involved. Their chief deficiency arises when a pilot is willing to give up speed for a quick position advantage. In this case, the aircraft with the P_S advantage will not necessarily win [19].

There have been several proposed extensions to the E-M format to improve its usefulness. The first discussed here is what Skow calls the extended E-M format [2]. The E-M plot of Figure 2.6 is simply stretched as in Figure 2.8 to include 0 g maneuvers and extremely high angles of attack. The 0 g condition is important because pilots can minimize their induced drag by unloading to 0 g. The P_S is therefore at a maximum here. Some aircraft can trim at angles of attack greater than that at maximum lift. It is useful then to plot the E-M characteristics at such extreme conditions because the maximum deceleration

potential occurs at maximum α. This extension gains even more importance as aircraft designs move into the post-stall regime.

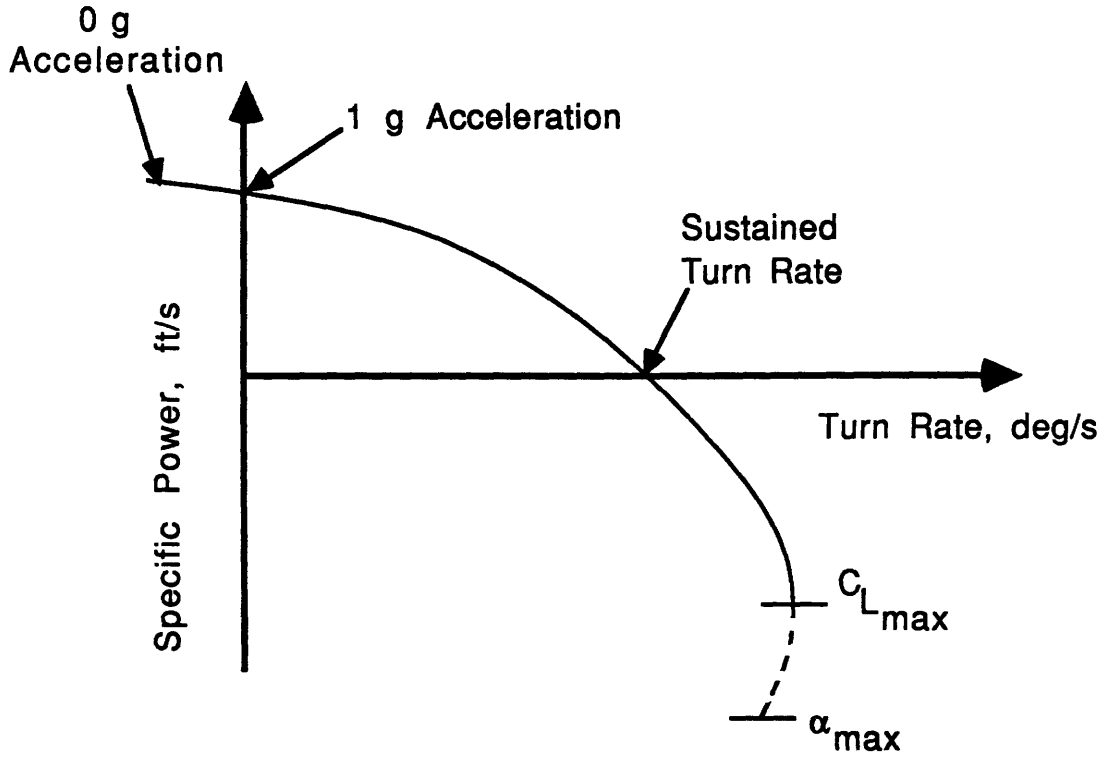

Figure 2.8. Extended E-M format. The maximum turn rate occurs at $C_{L_{max}}$, but the maximum deceleration potential is at α_{max} [2].

The second proposed modification to the E-M format is the dynamic speed turn (DST) plot, an extension of the doghouse plot [6,19,21]. The important parts of the doghouse are the maximum lift line, the maximum g line, and the 1 g line. Accordingly, the rate of change of airspeed is calculated for each point on the boundary of the doghouse. Two plots can then be constructed as shown in Figure 2.9.

Figure 2.9(a) contains points from the maximum lift and load factor lines. Using this plot, a pilot can calculate roughly his airspeed loss and average turn rate during a maneuver. In the example from [19], suppose a pilot begins a maximum maneuver at 500 kts. The estimated average deceleration is 20 kts/s, and the average turn rate is about 18 deg/s. Over the course of a 10 s maneuver, the aircraft will lose roughly 200 kts airspeed and will turn about 180 deg.

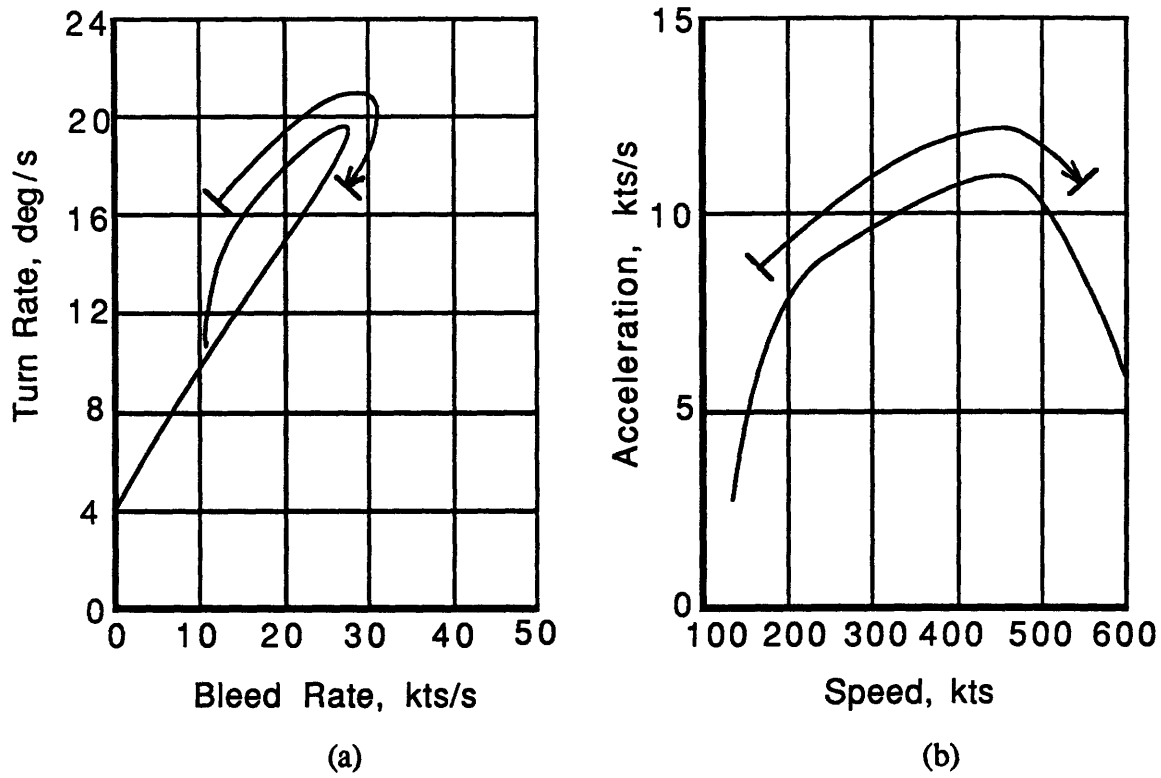

Figure 2.9. Dynamic speed turn plots. (a) shows turn rate and airspeed bleed rate for a maximum maneuver. (b) shows acceleration potential for a 1-g maneuver. The lines with arrows refer to examples discussed in the text [19].

Performance at 1 g is seen from Figure 2.9(b). In another example from [19], an aircraft initially at 200 kts has an average acceleration of approximately 10 kts/s over 30 s. The final airspeed is 500 kts.

The last E-M metric we will examine is also concerned with the time effects of combat maneuvering. In an m vs. n scenario, it is foolhardy to trade all of the energy advantage for a position advantage because once a kill is made there are other threats to worry about. The time necessary to regain the lost energy becomes critical. Dorn [4] proposes an energy-agility metric which accounts for both the time needed to kill and recover and the energy changes associated with the task. The metric is derived from the sample plot shown in Figure 2.10 and is defined as the shaded area:

$$\text{Energy} - \text{Agility} = \int_{t_o}^{t_{k/r}} \Delta h_e \, dt$$

where $t_{k/r}$ is the time to kill and recover and h_e is the specific energy.

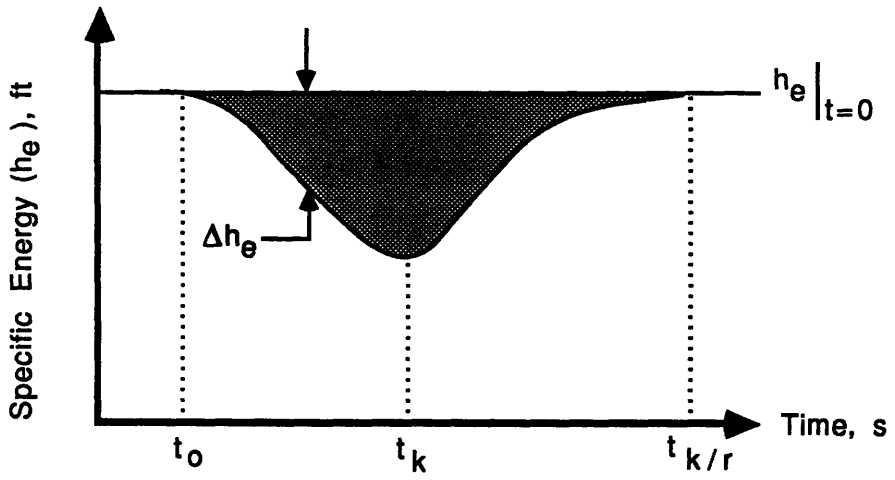

Figure 2.10. The energy-agility metric is the integrated product of specific energy and time [4]

There are ways to examine sustained performance other than through the E-M format. The DT parameter considers turn radius and turn rate simultaneously in evaluating agility. DT is defined as the product of crossrange needed to effect a horizontal turn through 180 deg and the time needed for such a turn [22,23]. An aircraft with a smaller DT achieves the first firing opportunity in a horizontal turn, that is, its t_k is smaller (Figure 2.11). The aircraft on the left has the smaller DT. Cannon [7] has used trajectory optimization to show aircraft on a minimum-DT trajectory can outperform the same aircraft on a minimum-time trajectory.

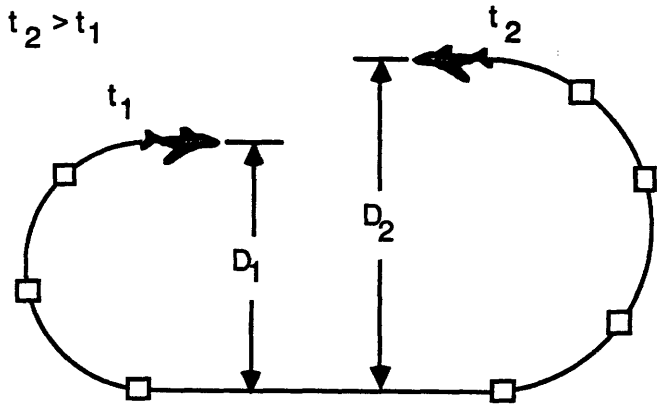

Figure 2.11. The DT parameter is the product of crossrange traveled in a 180 deg turn and the time to complete the turn. The aircraft on the left has a smaller DT and therefore the first firing opportunity. The squares are spaced an equal time apart [4].

Now we turn away from sustained performance to look at a performance measure used to quantify agility for shorter maneuvers. Torsional Agility accounts for an aircraft's ability to roll while loaded [5] and is defined as:

$$\text{Torsional Agility} = \dot{\psi} / t_{RC90}$$

where $\dot{\psi}$ is the aircraft turn rate and t_{RC90} is the time to roll and capture a 90 deg bank angle change at the given turn rate. Manned simulations [5] have suggested that differences in roll agility are apparent only after a few seconds. Thus Torsional Agility may be more appropriate for maneuvers of moderate duration than in point-and-shoot scenarios.

The final metrics to be examined focus more on instantaneous dynamics. They were proposed by Herbst [3] and approach agility mathematically rather than graphically. Second derivatives of the velocity vector are combined in various ways to form the agility vector. This vector can be resolved into three components: longitudinal, curvature, and torsional agility. It turns out these agility components are very similar to the time rates of change of N_x, N_z, and P_{stab}, respectively. An advantage to this approach is that the components are easily translatable to design quantities such as high pitch rate and roll accelerations and rapid engine spool times. These metrics help fill a void in agility research by quantifying transient agility, which is becoming more important as combat time scales shrink. For example, in the maneuver of Figure 2.12, much of the time is spent in transition rather than steady state.

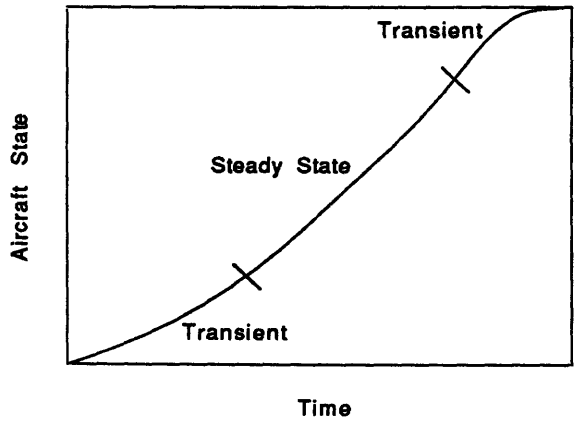

Figure 2.12. Transient dynamics form a large part of aircraft combat maneuvers [3]

3 Research Procedure

3.1 Plan

The object of the research was threefold. The first objective was to compare minimum-time and standard pointing maneuvers to see what time advantages could be realized merely by using optimal trajectories. The second objective was to see the effects of performance changes on aircraft agility by examining the new optimal trajectories and minimum times. The third goal was to demonstrate that agility comparisons based on a particular maneuver instead of optimal maneuvers may ignore important performance differences.

The pointing simulations began by imagining an aircraft to be in the center of a sphere, flying at the initial conditions given in Table 3.1. Next a matrix of 6 azimuth angle and 9 elevation angle pairs designating 54 targets was constructed. The azimuth angles were 15, 45, 90, 135, 165, and 179 deg; while the elevation angles were ±88, ±60, ±45, ±20, and 0 deg. The azimuth of 179 deg and the elevation of ±88 deg were chosen instead of 180 deg and ±90 deg to avoid singularities. The azimuth/elevation pairs define points only in the right half of the imaginary sphere in which the aircraft flies because symmetry was assumed. A few targets in the left hemisphere were tested to check this assumption.

In each of 54 test runs, the aircraft was made to acquire and stabilize on one of these attitudes in minimum time. Stabilization meant that each of the control variable rates was zero at the final time. This is not the same as pointing at a non-moving target because downrange and crossrange traveled was not considered. Rather, the maneuvers resemble the pitch angle capture tests run by the Air Force Flight Test Center and NASA for agility metrics data gathering [24].

The first 54 test runs were made using standard maneuvers taught to pilots as the best way to point the aircraft [25]. For the upper hemisphere, the standard maneuver is an unloaded roll in which the pilot banks to rotate the lift vector to the desired maneuver plane and then pulls up in that plane until the target is reached. The maneuver is called *unloaded* because the roll occurs at a load factor of one. For targets more than 15 deg below the horizon, the maneuver is the same, with the roll angle greater than 90 deg. This is because pilots cannot see more than about 15 deg below the cockpit. If a pilot is pursuing a target in this region, he would roll to see it before pointing the nose [25].

Table 3.1. Initial conditions

Quantity	Value
Mach number	0.9
Altitude	20000 ft
Velocity	933.24 ft/s
Weight	40700 lb
Flight path angle	0 deg
Azimuth	0 deg
Load factor	1 g
Angle of attack	2.14 deg
Angle of attack rate	0 deg/s
Roll angle about velocity vector	0 deg
Roll rate	0 deg/s
Thrust	45% of afterburner
Thrust rate	0 lb/sec

Because of the structure of the numerical optimization routine, it was easiest to model the standard maneuver as a minimum-time problem using limits on roll rate to allow rolling only at 1 g (see Section 3.3). Two alternate ways to formulate the problem are to specify a roll to the given maneuver plane and then a pull-up to the target in minimum time or to specify a roll angle and angle of attack history that the aircraft must follow regardless of time.

Once the 54 standard test runs were completed, the 54 corresponding optimal trajectories were found. Time histories of important state variables were kept so that differences from the standard maneuvers could be noted.

The next goal was to show the standard and optimal pointing times for any point on the sphere. This required a contour plot of pointing time as a function of attitude. It is easier to show time in a polar contour plot with maneuver plane (σ) as the angle coordinate and the total angle (TA) through which the nose sweeps as the radius. The maneuver plane and total angle are the first two rotations in a 1-2-3 Euler system. The azimuth/elevation pairs are the first two rotations in the conventional 3-2-1 Euler system. Maneuver plane and total angle could also be termed roll angle and pitch angle but are not to avoid confusion with the conventional Euler angles.

To get an accurate contour plot, many data points are needed. Unfortunately, only 108 were available, and it was not possible to perform more pointing trials because of the computational time involved. The problem was solved by using a standard multi-dimensional interpolation program [26]. Using this code, the pointing times for 550 attitudes were interpolated. There were 22 total angles in the interpolation matrix (from 15 deg to 172.5 deg in increments of 7.5 deg) and 25 maneuver planes (from -180 deg to +180 deg in increments of 15 deg). This assures complete coverage of the polar contour plot. Within the interpolation program, these attitudes were converted to azimuth/elevation form. Using the 108 test points, also in azimuth/elevation form, the interpolation was next performed. The attitudes were then converted back to σ/TA form, and the polar contour plots were constructed.

The σ/TA pairs and the corresponding times appear as in Figure 4.1. This plot is how the imaginary sphere looks to the pilot at its center, facing into the page. The dashed radial lines represent the maneuver plane, and the circular dashed lines represent the total angle through which the aircraft turns. For example, a pull-up maneuver follows the radial at 90 deg, and a 30 deg bank and total angle sweep of 150 deg lies at the intersection of the 60 deg radial and the third concentric circle.

The standard and optimal plots were overlaid to see the time advantages gained by performing the optimal maneuvers. The plots may also be used to see the maneuver planes in which aircraft are most agile. For example, the contours in Figure 4.1 indicate the aircraft can reach a point in the lower, forward quadrant quicker than one in the upper, rear quadrant.

As mentioned in the introduction, the nominal model had roll rate, angle-of-attack rate, and engine spool time limitations, discussed more fully in Section 3.3. Three new models were obtained by removing each of the three performance limits in turn. The trial/interpolation/plot sequence was repeated for each of the three models. The resulting contour plots were overlaid in turn with the optimal plot for the limited model to see how much the time improved. The state variable histories were also compared to each other and that of the limited model.

As will be discussed in Chapter 4, the results of these new simulations demonstrate that pitch rate limits significantly degrade performance. This fact was used to meet the third objective: show that comparing agility using standard maneuvers may mislead one to believe there are no agility differences in two quite different aircraft. The nominal model modified to have unlimited pitch rate was again used, this time in standard maneuvers. These standard runs of the nominal and unlimited models were compared on a polar plot.

Then the standard and optimal runs of the unlimited model were compared to show what performance capability the standard maneuver was concealing.

3.2 Aerodynamic Model

As was discussed in Chapter 2, only the force equations of motion are integrated in the simulations; the three moment equations are neglected. Roll angle and angle of attack are commanded directly, without regard to control surface deflections. This greatly simplifies computation because the tables needed to store aerodynamic data are smaller [27]. Lift and drag coefficients are functions of Mach number and angle of attack only, while thrust is a function of Mach number and altitude. These three forces, gravity, the equations of motion, and the commanded attitude variables completely define the aircraft position and velocity.

The aerodynamic model used was that of a high-performance fighter resembling the F-15 [27]. Lift, drag, and thrust as functions of flight condition were included in the model and are shown in Tables 3.2 - 3.4. OTIS fits a surface to these data, so it is important that they contain no sharp discontinuities. This is demonstrated in the carpet plots of Figures 3.1 - 3.3.

Weight, planform area, and specific fuel consumption were also needed. The values used were 40,700 lb, 608 ft^2, and 30,000 lb/hr. Constant specific fuel consumption was used for simplicity even though it changes with flight condition. The short maneuver times involved (less than 15 s) diminished the importance of this variable anyway.

Table 3.2. Coefficient of lift as a function of angle of attack and Mach number

Mach	Angle of Attack (deg)													
	-12	-8	-4	0	4	8	12	16	20	24	28	32	36	40
0.2	-0.79	-0.53	-0.27	-0.01	0.25	0.49	0.75	0.98	1.18	1.33	1.44	1.50	1.50	1.45
0.4	-0.84	-0.57	-0.30	-0.03	0.24	0.50	0.76	0.99	1.18	1.32	1.40	1.46	1.47	1.45
0.6	-0.89	-0.61	-0.33	-0.05	0.23	0.52	0.78	1.00	1.19	1.30	1.36	1.41	1.43	1.45
0.8	-0.98	-0.65	-0.36	-0.05	0.25	0.56	0.80	1.02	1.20	1.32	1.38	1.43	1.45	1.45
0.9	-1.07	-0.73	-0.38	-0.05	0.27	0.61	0.82	1.03	1.20	1.33	1.40	1.44	1.46	1.45
1.0	-1.10	-0.80	-0.43	-0.07	0.29	0.63	0.88	1.10	1.27	1.42	1.51	1.53	1.53	1.53
1.1	-1.00	-0.73	-0.39	-0.07	0.28	0.59	0.89	1.16	1.37	1.37	1.37	1.37	1.37	1.37
1.2	-0.90	-0.66	-0.34	-0.06	0.25	0.54	0.82	1.07	1.33	1.33	1.33	1.33	1.33	1.33
1.4	-0.80	-0.56	-0.28	-0.03	0.22	0.47	0.70	0.93	1.13	1.13	1.13	1.13	1.13	1.13
1.6	-0.70	-0.49	-0.24	-0.02	0.20	0.41	0.61	0.81	0.99	0.99	0.99	0.99	0.99	0.99

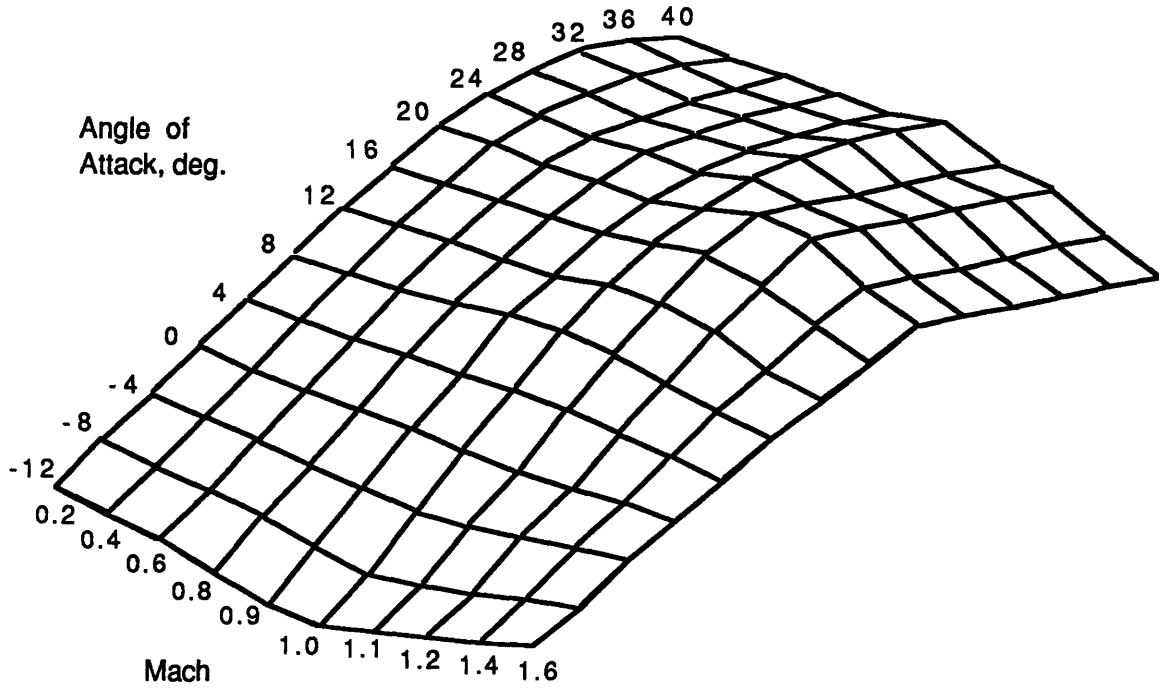

Figure 3.1. Coefficient of lift as a function of angle of attack and Mach number

Table 3.3. Coefficient of drag as a function of Mach number and coefficient of lift

C_L	Mach									
	0.2	0.4	0.6	0.8	0.9	1.0	1.1	1.2	1.4	1.6
0.0	0.0196	0.0196	0.0196	0.0201	0.0204	0.0261	0.0396	0.0466	0.0449	0.0430
0.2	0.0218	0.0218	0.0218	0.0218	0.0227	0.0321	0.0462	0.0523	0.0525	0.0525
0.4	0.0376	0.0376	0.0376	0.0370	0.0380	0.0554	0.0772	0.0860	0.0912	0.0967
0.6	0.0778	0.0778	0.0768	0.0745	0.0753	0.1037	0.1340	0.1496	0.1698	0.1906
0.8	0.1517	0.1517	0.1545	0.1583	0.1582	0.1809	0.2136	0.2420	0.3000	0.3600
1.0	0.2582	0.2582	0.2691	0.2862	0.2993	0.3166	0.3379	0.3711	0.4000	0.5000
1.2	0.4438	0.4438	0.4663	0.5092	0.5275	0.5370	0.5400	0.5500	0.5600	0.5800
1.4	0.8559	0.8559	0.8774	0.9000	0.9050	0.8787	1.0000	1.1000	1.2000	1.4000
1.6	1.2873	1.2873	1.3118	1.3222	1.3230	1.5000	1.6000	1.7000	1.8000	2.0000

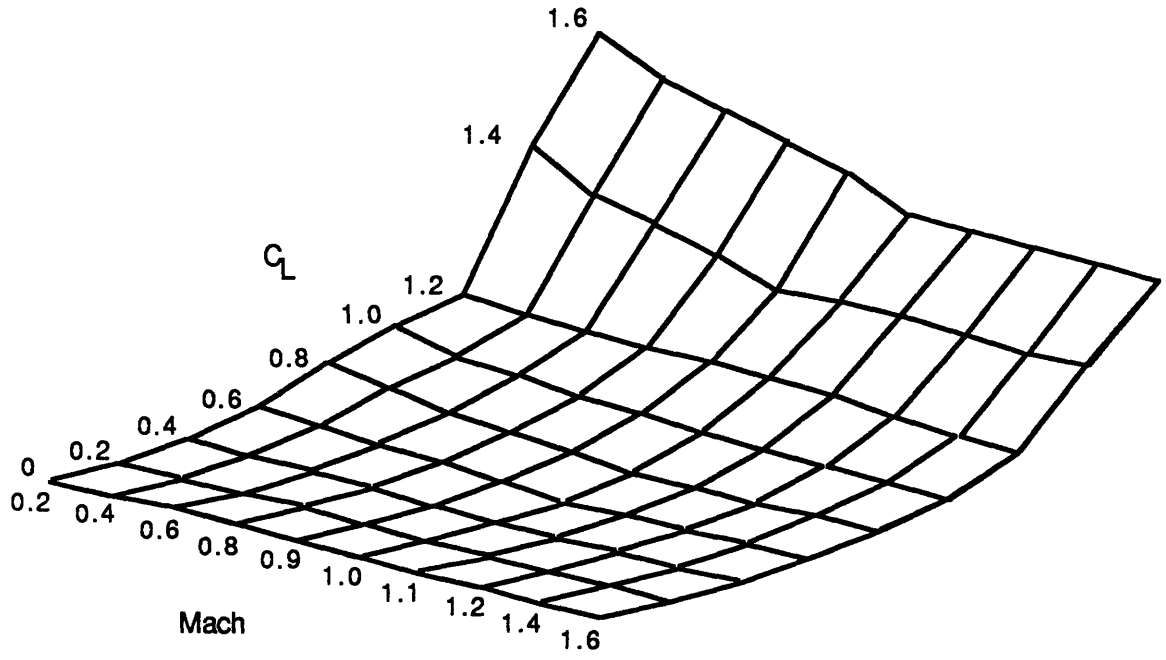

Figure 3.2. Coefficient of drag as a function of Mach number and coefficient of lift

Table 3.4. Afterburner thrust in pounds as a function of Mach number and altitude

Altitude	Mach									
(1000 ft)	0.0	0.2	0.4	0.6	0.8	0.9	1.0	1.2	1.4	1.6
0	18587	20022	21288	22814	24270	25355	27001	30219	35414	35414
10	13008	14190	15865	18258	20207	21283	22488	25103	29156	33724
20	9033	10145	11222	12804	15817	16631	17945	20083	23212	26734
30	6452	6452	7218	8532	10464	11863	13283	15657	17898	20569
35	4802	4802	5385	6401	7867	8921	10103	13055	15029	17277
40	3829	3829	4313	5144	6359	7211	8226	10606	12260	14179
45	3156	3156	3156	3835	4784	5456	6342	8045	9306	10944
50	2226	2226	2226	2770	3526	4062	4889	6039	6984	8319
55	1889	1889	1889	1889	2513	3032	3670	4386	5211	5261
60	1379	1379	1379	1379	1767	2167	2568	3068	3865	4672

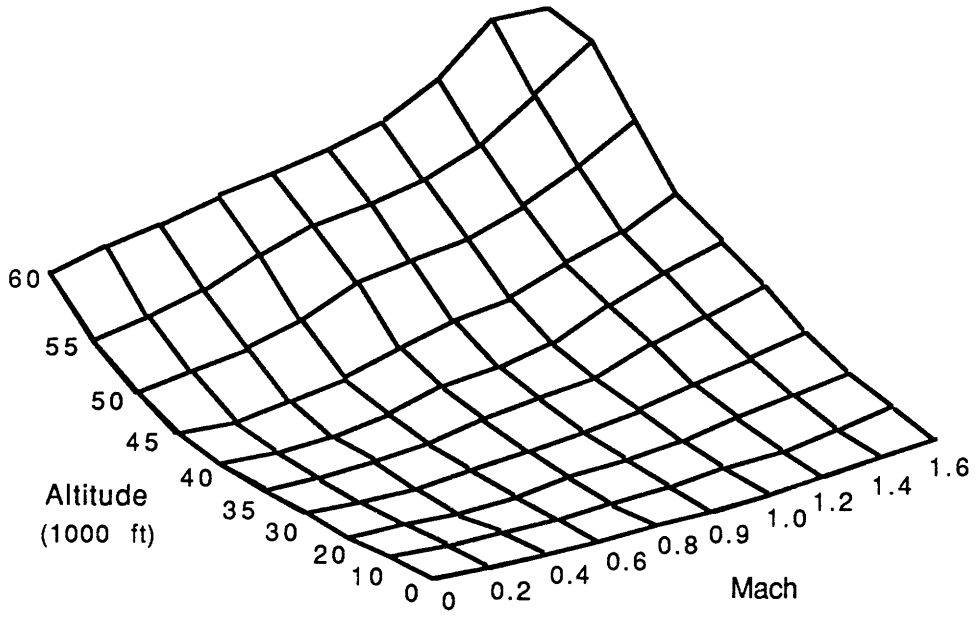

Figure 3.3. Thrust as a function of Mach number and altitude

3.3 Performance Limitations

Three performance limitations–roll rate, angle-of-attack rate, and engine spool time–were added to the model. Graphical representations of these limits are shown in Figures 3.4 - 3.6. The plane was also limited to an angle-of-attack range of -12 deg to 40 deg and a normal load factor range of -3 to +9. These values are representative of a modern fighter [25]. No lower bound was placed on airspeed.

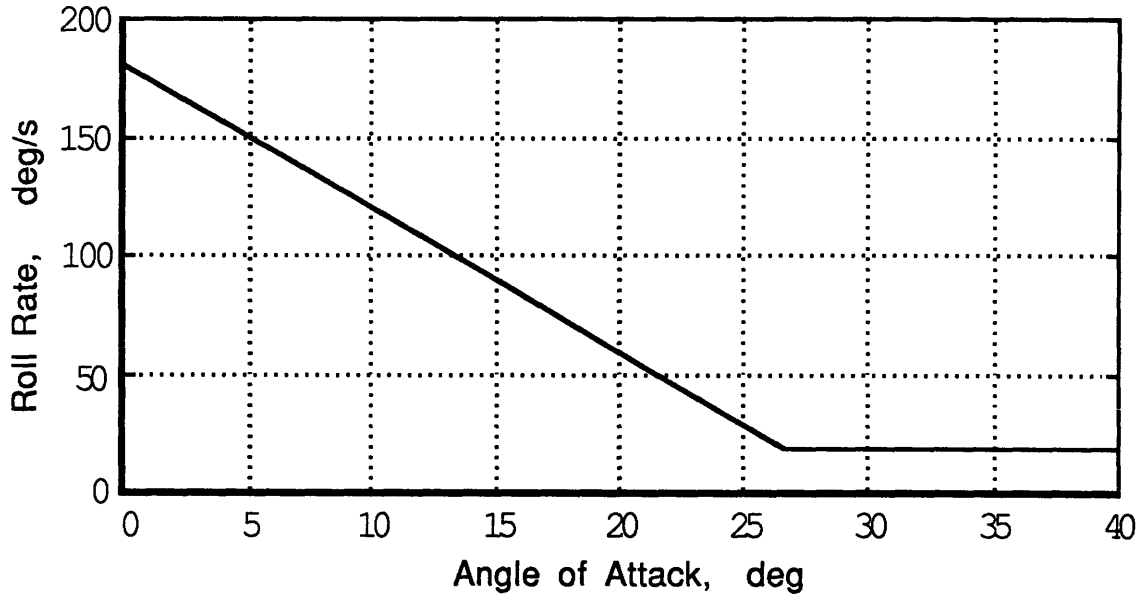

Figure 3.4. Roll rate limit as a function of angle of attack

Figure 3.4 shows that there is a significant reduction in roll rate at high angles of attack from the maximum of 180 deg/s available in straight and level flight. The three main reasons for this are reduced control surface effectiveness, adverse yaw, and roll coupling. Reduced control surface effectiveness simply means that the ailerons are less effective at the low speeds typically encountered at high angle of attack and may be completely useless after stall. There simply is not enough dynamic pressure to react against the ailerons. Adverse yaw occurs because of the aileron configuration during roll. If a plane is banking to the right, the left aileron deflects down, increasing left wing lift. The lift increase causes a corresponding increase in induced drag on the left side of the airplane, yawing the nose to the left and away from the bank direction [10]. The yawing motion may even cause flow past the vertical tail to induce a roll in the direction opposite the commanded roll. This is called roll reversal. Roll coupling is a result of adverse yaw. The commanded roll and the

adverse yaw combine to produce nose-up moments so strong that stall may result [2]. For this reason, fighters may have limiters in their control systems to prevent large roll rates at high angle of attack.

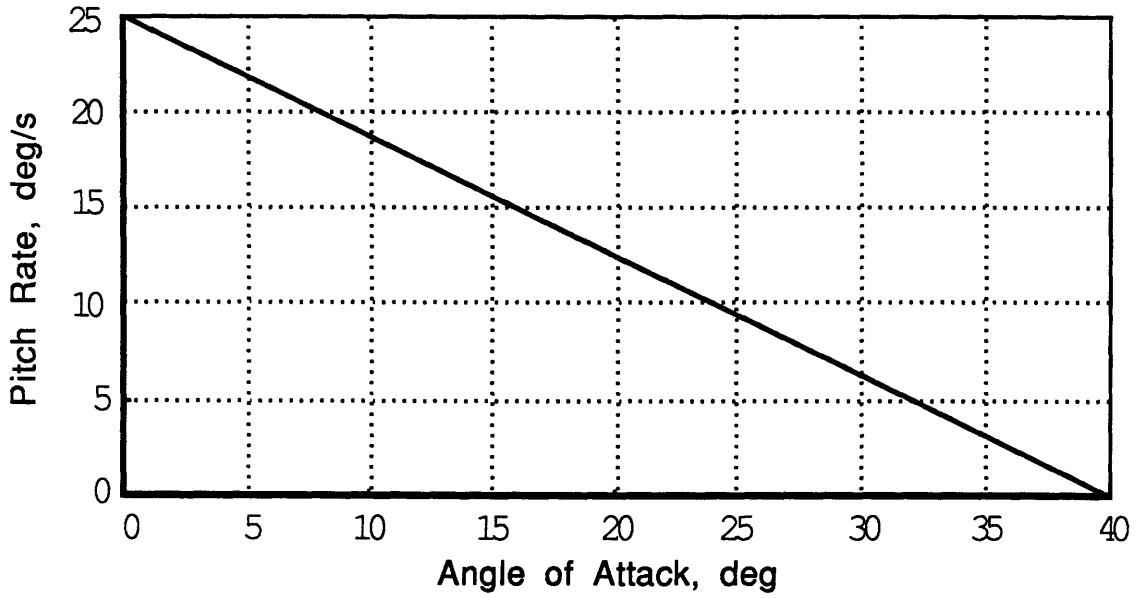

Figure 3.5. Angle of attack rate limit as a function of angle of attack

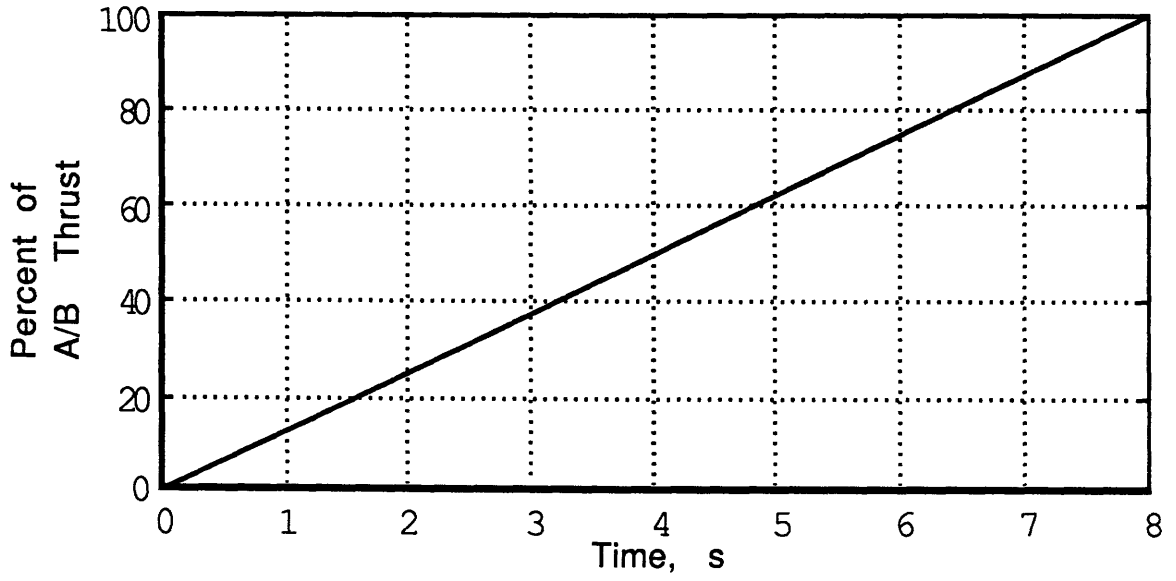

Figure 3.6. Thrust available as a function of time starting from idle at t=0

Angle-of-attack rate capability (Figure 3.5) decreases linearly from a maximum of 25 deg/s at 0 deg angle of attack to 0 deg/s at 40 deg. This is due to decreased control surface effectiveness and to longitudinal stability. For a longitudinally stable aircraft, increasing the angle of attack (and therefore the lift) increases the nose-down pitching moment, tending to return the aircraft to level flight. The effect is more pronounced at angles of attack near stall because the lift force is so great [10].

Pilots have three thrust settings from which to choose: idle, military power, and afterburner [25]. Military power is defined as the maximum thrust available with afterburners off and is about 45% of afterburner thrust at the initial flight condition of Table 3.1. Mainly because of turbine and compressor inertias, the thrust change between neighboring settings takes at least four seconds [28]. This means the pilot must wait eight seconds for the engine to spool up from idle to afterburner and eight more seconds to spool back down to idle. Thrust limits are shown in Figure 3.6.

The additional limit placed on roll rate for modeling standard maneuvers is shown in Figure 3.7. This limit constrains rolling to occur nominally at or below a 1.2 g load factor. The numerical optimization routine fits a cubic spline to these data, so it was necessary to add a gradual slope from the maximum roll rate at 1.2 g's to the minimum rate at 1.25 g's to eliminate spiking in the curve-fit. It was also necessary to allow a 10 deg/s roll rate at higher g's because the routine could not keep the roll rate at zero. That is not a severe weakness in the standard maneuver modeling: 10 deg/s is practically zero for a high-performance fighter.

The unlimited pitch rate model was also used to meet the third objective. For the standard maneuver, the load factor limit shown in Figure 3.7 was added simultaneously. The engine spool time and pitch rate limits were the same as previously discussed.

To help the reader keep all the models straight, Table 3.5 lists the limits used for each model, the types of runs performed with each model, and the name by which the model is referred in this thesis.

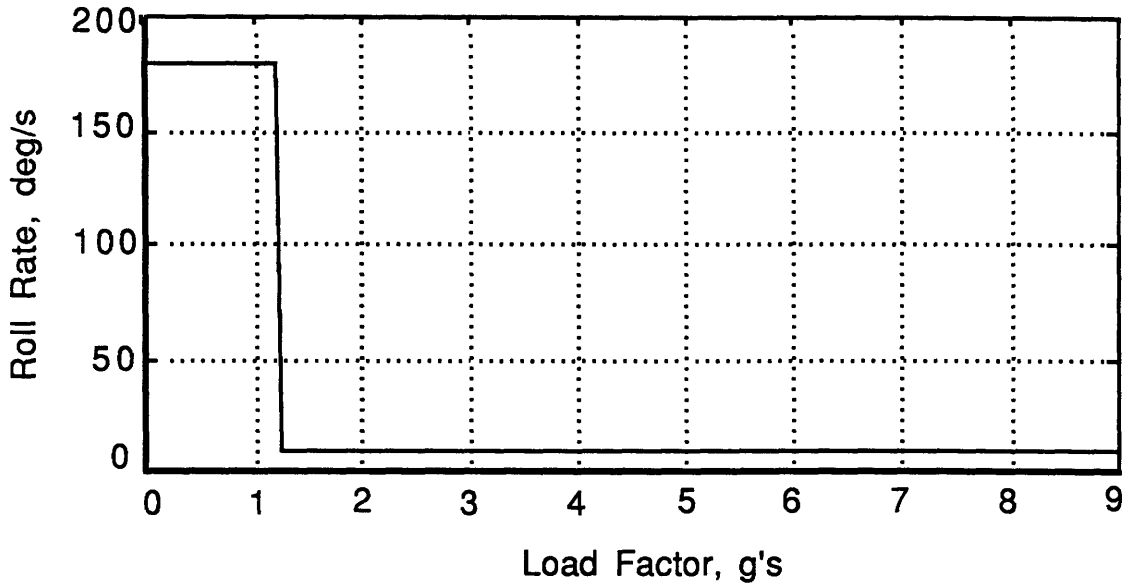

Figure 3.7. Limits placed on roll rate as a function of load factor to simulate standard maneuvers

Table 3.5. Model names, limits, and run types

Model Name	Figures Containing Model Performance Limits	Type of Runs Performed Using this Model	
		Standard	Optimal
Nominal	3.4, 3.5, 3.6, 3.7*	x	x
Unlimited roll rate	3.5, 3.6		x
Unlimited pitch rate	3.4, 3.6, 3.7*	x	x
Unlimited spool time	3.4, 3.5		x

* Used only for standard maneuvers

4 Results

4.1 Introduction

Chapter 3 discussed the research procedure in depth. Four models, one nominal and three unlimited, were to be tested using optimal and standard pointing maneuvers to 54 targets. The time histories would reveal the advantages gained by switching to optimal trajectories and the relative importance of different performance limitations. This chapter provides the simulation results in the form of tables of the pointing times, polar contour plots, and time histories of selected runs. Obviously not all 54 runs for each model could be included in a finite document, so runs representing targets in different locations about the imaginary sphere were chosen to give the reader a feel for the aircraft dynamics.

Some points concerning the time history plots need to be mentioned. Each state variable is scaled consistently among the plots to allow easier comparisons. For example, the drag coefficient scale always runs from 0 to 1.5. The first exception to this is altitude which is auto-scaled in each plot according to the minimum and maximum altitudes shown. The second exception is that the load factor scale in the optimal vs. standard plots runs from -5 to 10 instead of 0 to 10 as in all the other runs. The independent variable time is auto-scaled in each plot because of the wide time difference among runs. Acceleration, shown in the seventh plot of each time history set, refers to the time rate of change of speed at the center of gravity.

For brevity, targets are written as (135,45), where the first number stands for the azimuth and the second number stands for the elevation.

Figure 4.1 shows how a typical polar contour plot appears. The optimal trajectories of the nominal model are depicted in this figure. The plot is a two-dimensional mapping of the imaginary sphere in which the aircraft is flying. The airplane is flying into the paper at the center of the dotted concentric circles. The sphere is then unfolded such that the point directly behind the aircraft in the sphere is mapped into a concentric circle at 180 deg. These concentric circles represent the total angle the nose sweeps through to reach a target. The radial lines are maneuver plane or bank angle, the vertical radial representing straight up or down. In all the plots to come, contours of 2, 4, 6, and 8 seconds are drawn to show how far the aircraft can point in these specific times.

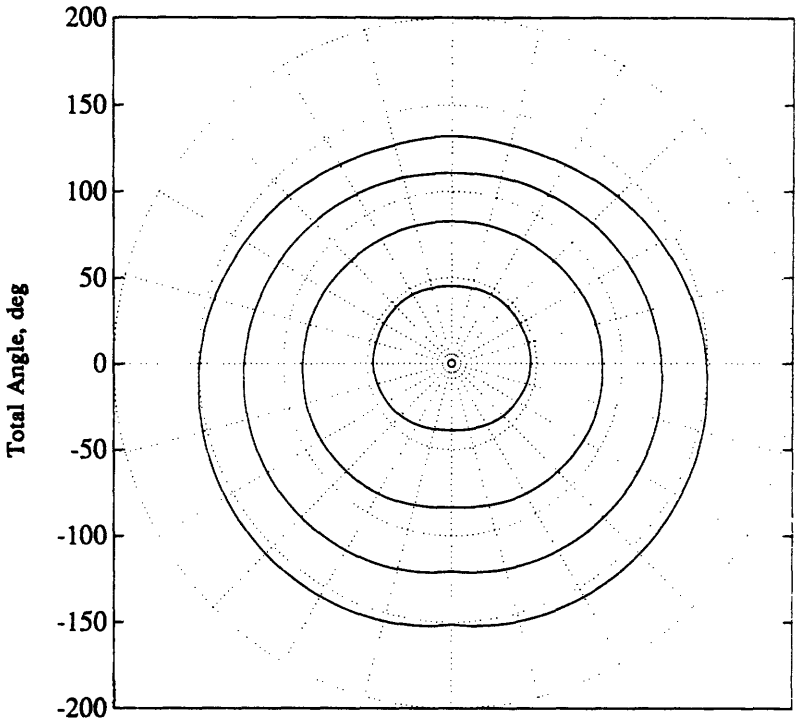

Figure 4.1. Typical contour plot. The solid lines represent time contours of 2, 4, 6, and 8 s. The curves become more egg-shaped as time progresses because of gravity. The slight bump at the top of the 8 s contour represents the switch to a split-S from a bank-and-pull.

Notice that the 2 s contour is nearly circular, while the 8 s contour is more egg-shaped. This is due to gravity. The longer the trajectory, the more important gravity becomes, either as a help or a hindrance. With the plot in Figure 4.1 as an example, we see that in 8 s the aircraft can travel about 20 deg farther straight down than it can straight up. Later contour plots are not so neatly explained because of the various limits on the models, but gravity always helps in lower hemisphere pointing.

In the 8 s contour there is a very slight bump near the 0 deg maneuver plane. It is more noticeable in other plots, especially Figure 4.2. This area of the plot represents heading angles of about 180 deg and elevation angles of 45 to 90 deg. Normally for rear hemisphere pointing the aircraft banks to the appropriate plane and pulls within the plane until the target is reached. But for some targets near the bump it is quicker to pull straight up and over in a split-S. The bank-and-pull maneuver in these cases is a local minimum, but the split-S is the global minimum.

The final general trend to mention is the slight cusps in the contours at the 180 deg maneuver plane. Recall that the interpolation program interpolated from the -180 deg to +180 deg maneuver planes. Thus the heading angle at the cusps falls right at the edge of the interpolation, and the algorithm does not produce smooth results at this point.

4.2 Agility Differences due to Trajectory Optimization

4.2.1 Time matrices

Tables 4.1 and 4.2 contain the final times for each of the 54 targets tested. In all cases save one, the optimal maneuver times are less than or equal to the standard maneuver times. The exception is for (135,88), in which the optimal time is 0.01 s slower than the standard time. This is attributable to finite tolerances in the digital computer performing the simulations. Practically, the two times are identical. The most significant time difference was for the target (15,0). The optimal time is 48% faster than the standard time.

Table 4.1. Pointing times for optimal trajectories of nominal model

Heading Angle (deg)	Pitch Angle (deg)								
	-88	-60	-45	-20	0	20	45	60	88
15	4.21	3.00	2.35	1.29	0.81	0.96	2.07	2.77	4.33
45	4.25	3.26	2.76	2.16	1.93	2.10	2.67	3.19	4.36
90	4.31	4.13	4.08	4.02	4.08	4.14	4.31	4.38	4.46
135	4.39	5.22	5.68	6.34	6.85	7.06	6.50	5.84	4.56
165	4.39	5.78	6.58	7.92	9.10	9.84	8.19	6.68	4.59
179	4.41	5.95	6.93	8.48	9.76	11.00	8.35	6.77	4.60

Table 4.2. Pointing times for standard trajectories of nominal model

Heading Angle (deg)	Pitch Angle (deg)								
	-88	-60	-45	-20	0	20	45	60	88
15	5.01	4.29	3.41	1.69	1.56	1.18	2.21	2.77	4.34
45	4.96	4.27	3.89	3.13	2.54	2.43	3.16	3.81	4.36
90	4.92	4.66	4.49	4.52	4.67	4.81	5.04	5.16	4.46
135	4.94	5.55	6.13	7.07	7.67	8.01	7.61	6.95	4.55
165	4.98	6.16	7.21	8.69	9.90	10.89	8.35	6.84	4.59
179	5.01	6.39	7.65	9.46	10.85	12.13	8.35	6.77	4.60

4.2.2 Polar contour plots

Figure 4.2 contains the polar contour plot comparing the two types of maneuvers. Almost everywhere the optimal maneuver is superior to the standard maneuver. The strong exception occurs along the vertical radial in the upper hemisphere where the two exhibit the same performance. This region represents a near pull-up with no bank angle. Recall that the standard maneuver constrains the aircraft to roll at 1 g. Because no bank angle is required for these targets, there is no penalty on the standard maneuver.

This figure clearly shows the effect of local vs. global maxima. The pronounced bumps near the upper hemisphere vertical radial come from the transition to split-S maneuvers from bank-and-pull maneuvers. The bumps are more noticeable than in Figure 4.1 because a bank-and-pull maneuver at 1 g is much slower than a loaded roll. The time difference between a split-S and a bank-and-pull is therefore magnified.

4.2.3 Time Histories

The time histories in Figures 4.3 through 4.8 show that a loaded roll is much more effective than a roll at 1 g for rapid pointing. In all of the plots except Figure 4.3, the optimal maneuver is to roll and pull simultaneously. In Figure 4.3, the bank angle for (179,45) is zero, so there is no difference between the optimal and standard trajectories.

The maneuver for pointing to (179,45) is a split-S. The aircraft pulls up until it reaches an angle of attack of about 18 deg. This corresponds to the load factor limit of 9 g's. The load factor depends most heavily on speed and angle of attack. Thus the airplane waits at 18 deg until enough airspeed is bled off that it can continue pulling without violating the 9 g constraint. After the aircraft begins falling from the g_{max} limit, α is free to increase again and does do until it reaches the imposed limit of 40 deg.

Thrust increases almost immediately until A/B. The time delay of 4 s from 50% thrust to 100% thrust clearly appears in the thrust history.

There are two observations from Figure 4.3 that hold for all the time histories in this chapter. The altitude depends on which way the lift vector is oriented and not very much on speed. In this example, the lift vector is oriented up from the pilot's perspective, so the aircraft climbs. This is consistent with the dropping speed, but in some future cases we will see a drop in both altitude and speed. The second observation concerns the drag coefficient history: it closely follows angle of attack. The angle of attack plateau mentioned was seen in most other time histories, too, but there are additional causes for it depending on the particular model.

Figure 4.4 is for the target (135,45) in the upper right quadrant of the polar contour plot. The transformation to maneuver plane/total angle yields a 35 deg bank followed by a pull through 120 deg. However, the figure exhibits an increasing bank angle because σ is the bank about the velocity vector, not the aircraft longitudinal axis.

We see here a dramatic difference in the load factor history between standard and optimal trajectories and a correspondingly large difference in the final times. The bank angle histories are nearly the same, but the standard maneuver is to wait until the bank is essentially completed before beginning the pull. The optimal maneuver is a simultaneous roll and pull. Consequently the shapes of each plot are the same with the optimal plots shifted back in time about 1 s.

Again the thrust increases immediately to A/B with the spool time limits in effect. The speed loss slows after the thrust increase is well underway. The angle of attack plateau corresponds roughly with the acquisition of g_{max}, but α soon starts up again and eventually reaches the α_{max} of 40 deg. Drag closely follows α.

Figure 4.5 compares the two trajectories to (135,45) in a three-dimensional manner. The airplanes are drawn every 0.8 s. The path taken by the center of gravity is traced and projected onto the ground plane. The airplanes in a given trajectory become closer as time progresses, demonstrating the slow-down as angle-of-attack increases. The optimal maneuver is to bank and pitch simultaneously, pulling the aircraft center of gravity in the bank direction almost immediately. The standard trajectory center of gravity tracks straight for about 2 s, corresponding to the pitch onset. As lift increases ,the center of gravity is pulled right. The standard downrange is about 1000 ft more than optimal. The crossranges traveled and the final altitudes are almost identical.

Figure 4.6 contains histories for a horizontal turn to a heading of 90 deg. Again the delay in angle of attack until the appropriate bank angle is reached is evident in the standard trajectory, as is the time shift for the optimal trajectory.

As before, the thrust history is not shifted. New for this target, though, is the noticeable difference in the thrust histories. Also in contrast to the previous figures, neither thrust reaches A/B. We will see in Figure 4.13 that with no spool time limits, the thrust will throttle down to idle and then up to A/B. However, with the imposed spool-time limits, this is impossible. Instead what happens is that the thrust stays at the average of idle and A/B, or nearly at military power. The two thrusts increase at approximately the same slope. The standard maneuver is longer, so the final thrust is higher.

The angle of attack reaches a maximum of about 30 deg, well below the maximum allowable of 40 deg. The reason is the limit placed on pitch rate. This limit is a linear decrease from 25 deg/s at $\alpha = 0$ deg to 0 deg/s at $\alpha = 40$ deg. Accordingly, we can see the slope lessen as α increases. The slope flattens out completely at the end, a consequence of the restriction that control rates be zero at the end of the trajectory.

The airplane climbs slightly in the optimal trajectory because there is no waiting for the lift vector to orient itself before the angle of attack rise begins.

We now look at a lower hemisphere target, as shown in Figure 4.7. This depicts a point to a heading angle of 135 deg and an elevation of -20 deg. We see all the same phenomena discussed earlier except the thrust behavior. This time the optimal thrust is noticeably higher than the standard thrust. This is because the immediate rise in α for the optimal trajectory causes the speed to drop faster. The thrust goes towards A/B to keep the speed up and the airplane at g_{max} as long as possible.

The last comparison in this section is between the standard and optimal trajectories to (15,-88), a target almost directly beneath the initial position, and is shown in Figure 4.8. Again the optimal α history is shifted. The standard α history dips to about -5 deg before climbing. Remember that no penalty was placed on rolling at negative load factors or negative angles of attack (Figure 3.7). A negative α with a dropping altitude means the nose is starting to point downwards. This allows the aircraft to begin the point to the elevation of -88 deg while still rolling. Remaining at the initial α while rolling would eliminate this headstart. (This angle of attack dip was the reason the standard vs. optimal plot load factor scale begins at -5 instead of 0.)

The optimal bank angle more closely approximates a step than the standard bank angle does. The standard bank angle increases gradually after the large initial jump to 150 deg. This figure confirms that modern pilots adopt the correct strategy of rolling 180 deg for lower hemisphere targets. Not only does this help the pilot see threats beneath him, it improves his pointing time.

Both thrust histories slope downward initially. Then the optimal thrust starts back upward, while the standard thrust continues towards idle without reaching it. The optimal thrust must increase because the much higher drag causes a large velocity drop. Speed

must be high to stay on the g_{max} line. On the other hand, standard thrust drops because the drag is so low that speed remains constant for quite awhile. Were thrust and speed to increase, the load factor constraint would have been violated as α climbed.

This example points out how crucial drag is in slowing the airplane. Even though the airplane is pointing nearly straight down, the total speed loss is about 300 ft/s. This is because the high angles of attack contribute to the total drag by greatly increasing the induced drag. Notice that the minimum C_D occurs at a load factor of 0 g. It was mentioned in Section 2.2 that pilots often unload to 0 g to minimize induced drag and maximize specific power.

The last set of plots in this section, Figure 4.9, was included to show how points close together in the imaginary sphere might be reached by quite different optimal paths. The figure depicts the targets (179,45) and (179,20) for the nominal optimal model. These two targets are on the upper half of the vertical radial, separated by a total angle of only 25 deg. Yet there is a big difference in the optimal way to reach each point. (179,45) requires a split-S: the bank angle is zero. The aircraft banks and pulls to reach (179,20): the bank angle is not zero.

4.3 Agility Differences due to Changes in Performance Limitations

4.3.1 Differences due to Removing Spool Time Limits

4.3.1.1 Time matrix

Table 4.3 shows that the unlimited spool time model acquires targets faster or in at least the same time as the nominal model of Table 4.1. The margin is very small, however. The largest time difference is 0.2 s for the heading angle/pitch angle pairs of (179,20) and (165,45).

4.3.1.2 Polar contour plots

Figure 4.10 confirms that there is not much advantage gained by removing spool-time limits, at least for minimum-time pointing. The biggest time differential comes in the 8 s contour, with the differences slowly disappearing for close-in contours. This suggests spool time may be more important for longer trajectories in which energy management becomes critical.

Table 4.3. Pointing times for optimal trajectories of unlimited spool time model

Heading Angle (deg)	Pitch Angle (deg)								
	-88	-60	-45	-20	0	20	45	60	88
15	4.14	2.94	2.34	1.29	0.77	0.96	2.04	2.77	4.28
45	4.16	3.21	2.74	2.16	1.92	2.05	2.66	3.15	4.32
90	4.22	4.06	4.00	3.99	4.02	4.11	4.26	4.34	4.40
135	4.28	5.17	5.66	6.31	6.72	6.91	6.34	5.76	4.49
165	4.31	5.72	6.56	7.82	8.97	9.66	7.99	6.59	4.53
179	4.32	5.77	6.90	8.38	9.65	10.80	8.16	6.59	4.55

4.3.1.3 Time histories

The thrust histories can vary greatly depending on the target. Figures 4.11 and 4.12 reveal that the optimal thrust strategy for targets at θ = -90 deg is to pull to idle immediately and then switch to A/B just before the load factor starts slipping from the maximum of 9 g's. The load factor is dependent chiefly on angle of attack and speed. The pull to idle slows the aircraft down to allow higher velocity vector turn rates (as corner speed is approached) and to allow more angle of attack without violating load factor limits. High angles of attack are important to allow the nose to point without changing the velocity vector much; but they also create very high drag, slow the aircraft, and cause a drop from the maximum load factor. This speed penalty is why the maximum angle of attack of 40 deg is not reached in these two cases.

Even with the greater spooling capability, there is not much of an effect on velocity when the thrust is pulled back to idle (see the velocity portions of Figures 4.11 and 4.12). This is because the tremendous inertia present at the initial condition of Mach 0.9. It takes a long time to slow down an aircraft going this fast, but there is enough of a time savings to warrant throttling between idle and A/B instead of just staying at military power as the nominal model has done.

Similar results hold for trajectories of moderate duration such as that shown in Figure 4.13. The maneuver time is similar (about 4 s), the angle of attack does not reach maximum, and the load factor remains at the maximum for a fairly long time. The benefits of the idle-to-A/B excursion are again slight but are present nonetheless.

Targets in the forward hemisphere generally do not warrant such dramatic changes in thrust. Figures 4.14 and 4.15 are good examples. In Figure 4.14 the thrust is near idle for most of the trajectory and does not return to military power, much less full A/B. The thrust does approach military power near the end of the trajectory in Figure 4.15, but is still not the bang-bang control seen earlier. The reason can be seen in the load factor, velocity, and drag coefficient plots of the figures. Compared to Figures 4.11 and 4.12, there is less drag and a correspondingly smaller drop in velocity. Hence the load factor stays at maximum for most of the trajectory. There is not a need to throttle to A/B to keep at g_{max}. The maneuver times are much lower, so the angle of attack is not high for very long. The drag coefficient cannot build to the levels seen in the previous cases.

For considerably longer trajectories like that shown in Figure 4.16, the angle of attack must be at or close to the maximum to allow faster pointing. This means drag can be significant over a long portion of the flight. Velocity drops rapidly, in this case to below 400 ft/s; and there is no danger of violating load factor constraints. In fact it is difficult to remain at g_{max}. Thus to minimize the pointing time, it is necessary to throttle immediately to A/B to keep the velocity as high as possible. The velocity loss is indeed less for the unlimited spool time model.

One of the purposes of removing the spool time limits was to see how fast the thrust would change given infinite capability. The fastest thrust change observed was 142% max thrust/s for the target (90,-45). This translates to a spool time of 0.7 s from idle to A/B.

4.3.2 Differences due to Removing Pitch Rate Limits

4.3.2.1 Time matrix

Table 4.4 shows that removing angle of attack rate limits consistently produces lower pointing times (compare Table 4.1).

4.3.2.2 Polar contour plots

The contour plot of Figure 4.17 tells us the most significant advantages come for total angles less than 100 deg. For example at 4 s, the unlimited model travels about 10 deg farther regardless of maneuver plane. There is practically no difference in the contours at 8 s. This lends credence to the supposition of Section 4.3.1 that for long trajectories energy management becomes more important. For such maneuvers, the model is not taking all the pitch rate capability available because high angles of attack cause high drag.

Table 4.4. Pointing times for optimal trajectories of unlimited pitch rate model

Heading Angle (deg)	Pitch Angle (deg)								
	-88	-60	-45	-20	0	20	45	60	88
15	3.65	2.78	2.14	1.09	0.57	0.58	1.73	2.19	3.59
45	3.91	2.94	2.67	2.07	1.80	1.80	2.17	2.36	3.64
90	3.67	3.53	3.49	3.43	3.43	3.47	3.60	3.68	3.76
135	3.78	4.74	5.34	6.25	6.80	7.01	6.43	5.66	3.89
165	3.74	5.42	6.39	7.88	9.03	9.82	8.13	6.61	3.94
179	3.76	5.35	6.76	8.40	9.66	10.95	8.29	6.70	3.94

4.3.2.3 Time histories

Four different groups of targets will be discussed in this section. The first group consists of three horizontal turns of increasing total angle, and the second group consists of three targets of differing total angle in the lower right quadrant of the contour plot, Figure 4.17. One target nearly straight down is examined in the third group. Group 4 contains one target in the upper right quadrant of the contour plot. The first two groups were chosen to explain how the differences between the nominal and unlimited model change as the total angle increases. The last two groups were included to demonstrate how trajectories in other areas of the contour plot look.

There are some observations which apply to all the plots. Often the unlimited model was able to reach α_{max} when the nominal model could not. If both models reached α_{max}, the unlimited model acquired it first. The angle of attack plateau at about 18 deg was achieved sooner and/or held longer for the unlimited model. Consequently g_{max} was also reached sooner and held longer. The load factor at the end of the unlimited model trajectories was higher.

Figures 4.18, 4.19, and 4.20 show how the trajectories change as the total angle increases from 45 deg to 90 deg to 135 deg in a horizontal turn. In all instances, the unlimited pitch rate model pops up to an angle of attack of roughly 18 deg, remains nearly constant, and then pops up again at the end of the trajectory. We see that for (45,0) and (90,0) the angle of attack plateau nearly coincides with the load factor limit of 9 g's. This means any sharp increase in α will violate the load factor limit. A gradual increase is allowed because the speed is dropping: a speed loss with no accompanying α rise would lower the load factor.

There is not such a clear correspondence between angle of attack and load factor for the (135,0) case. We turn to the drag history for an explanation. Between 2 and 3.5 s the speed is about 800 ft/s. At a 20,000 ft altitude, this gives a Mach number of just under 0.8. Table 3.2 shows that with the angle of attack of just under 20 deg in the same time interval, the corresponding C_L is 1.2. Therefore from Figure 3.2 we see a sharp increase in drag at the given flight condition. Because drag is strongly tied to angle of attack, the airplane does not continue its initial α rise to avoid this drag jump at $C_L = 1.2$ and Mach = 0.8. An increase in drag would drop the speed and hence the load factor. Figure 4.20 shows C_D is indeed lower for the unlimited model most of the time. The load factor history shows the nominal model has a higher load factor for about 1 s because of the higher α, but soon the higher speed loss overwhelms any advantage of the higher α.

The (45,0) trajectory is so short that α_{max} is never reached. If it were, the load factor limit would be violated. For (90,0) and (135,0), however, the maneuver is long enough to bleed off the airspeed necessary to allow the full α capability to be used. Only for (135,0) does the nominal model also reach α_{max}. This is because of the speed and angle of attack tradeoff for load factor. By the time the angle of attack of 40 deg is reached in the (90,0) case, the maneuver is essentially completed; so the resultant loss in airspeed does not impact the load factor history. The quick pop-up to capture the target is more important than any detrimental effects on speed. But for (135,0) energy must be more carefully managed. The time history differences are a result of this management.

Both the nominal and unlimited model thrust histories in Figure 4.20 confirm that for long trajectories, the optimal thrust strategy is to reach A/B as soon as possible. There is a 4 s delay from military power to A/B imposed by the engine inertia that is clearly evident from the thrust history.

We now look at targets in the lower right quadrant. The histories of interest are in Figures 4.21, 4.22, and 4.23. The first detail of note is the different angle of attack strategy compared to the horizontal turns. In that case, the unlimited model pulled up faster than the nominal model. With the new figures, however, the reverse is true. This is because the target is in the lower hemisphere. If the aircraft began pulling up immediately before the roll through 90 deg was completed, the elevation angle would increase instead of decrease. The nominal model pulls up sooner because its pitch rate is limited to 25 deg/s, and it needs the extra time to reach the appropriate angle of attack. The bank angle for the unlimited model leads the nominal model bank slightly because the sooner the bank angle reaches the desired value, the sooner the extra pitch capability can be used.

In Figure 4.21 there is no sudden pop-up at the end of the trajectory. This is because the load factor is at its limit until the end, so an α excursion would violate the limits. The effect of waiting until the bank angle is completed before initiating the pitch is

seen in the altitude plot: the nominal model climbs slightly away from the target elevation before dropping.

The thrust differences are evident in Figure 4.22: the unlimited model thrust decreases slightly, while the nominal model thrust increases. This difference is again due to the angle of attack and speed tradeoff for load factor. Because the nominal trajectory is longer, speed has more time to bleed off. In an attempt to keep speed as high as possible, thrust kicks in as soon as the speed drops enough that no violation of load factor limits can occur. This thrust increase can only occur as fast as the engine spool-up time allows. This time is 4 s from military power to A/B, and Figure 4.22 shows that in the 2 s from throttle-in, the thrust is about halfway to A/B.

This figure also demonstrates the quick α pop-up at the end of the maneuver for the unlimited model.

Figure 4.23 is a target with approximately the same maneuver plane as the previous two but a much higher total angle. Thus it is important to conserve energy by keeping the angle of attack and hence drag as low as possible for as long as possible. The figure shows that the drag coefficient for the unlimited model is lower than that for the nominal model over most of the trajectory. As before the thrust tries to hit A/B immediately subject to spool-time limits.

Figure 4.24 shows the history for (15,-88), pointing nearly straight down. The trajectory is like (90,-45) in most ways: there is a period of constant angle of attack as bank angle increases, an α plateau nearly coincident with the attainment of g_{max}, a pop-up at the end of the trajectory to capture the target, and a significant time advantage for the unlimited pitch rate model. The biggest difference is in thrust. The unlimited thrust for (15,-88) reaches idle, while the (90,-45) thrust remains nearly constant. The reason is that the additional altitude loss (about 400 ft) for (15,-88) is converted to kinetic energy that helps keep the speed up. Less thrust is therefore needed.

The last group concerns a target in the upper right quadrant of the contour plot, (15,45), shown in Figure 4.25. There is nothing special about this plot except for the much higher initial angle of attack for the unlimited model. This is because the trajectory is so short that the significantly higher drag does not have time to degrade the speed much.

This target was also chosen to demonstrate what happens when the trajectory is divided into finer segments. Recall that OTIS allows the user to choose how many segments will comprise the maneuver, and that 9 equal segments were used for this research. The trajectory might change significantly if more segments were used. There might, for example, be more step-like behavior in the unlimited model. Accordingly, the (15,45) maneuver was divided into 19 segments, with the first five segments being three times smaller than the last 14. This means the first 0.15 s was composed of five segments.

The resulting histories are shown in Figure 4.26. Compared with Figure 4.25 there are indeed important changes. The angle of attack history approaches a step and the load factor history follows suit. The α pop-up at the end remains. Banking begins slightly faster. The time for the unlimited model trajectory shortened from 1.73 s to 1.44 s. The nominal model completed the run in 1.85 s, down from 2.07 s. The Δt_f between the unlimited and nominal models climbed to 0.41 s from 0.34 s. Various other targets were also tested and the results summarized in Table 4.5. In all cases the times were significantly reduced, but the Δt_f's often were not. Because we are primarily interested in improvements resulting from different models rather than the absolute times, the Δt_f's are more important. It was this fact, coupled with the fact that runs using 19 unequal segments took over four times more CPU to complete, that made the nine-segment choice a reasonable one.

Table 4.5. Time differences between runs using nine equal and 19 unequal segments

Target	Nine Segments			Nineteen Segments		
	Nominal t_f	Unlimited t_f	Δt_f	Nominal t_f	Unlimited t_f	Δt_f
(135,0)	6.85	6.80	0.05	6.33	6.27	0.06
(165,45)	8.19	8.13	0.06	7.57	7.42	0.15
(179,60)	6.77	6.70	0.07	6.33	6.15	0.18
(90,0)	4.08	3.43	0.55	3.67	3.13	0.54
(15,45)	2.07	1.73	0.34	1.85	1.44	0.41

The maximum pitch rate ever used was 57 deg/s. The maximum allowable for the nominal model was 20 deg/s.

4.3.3 Differences due to Removing Roll Rate Limits

4.3.3.1 Time matrix

Table 4.6 shows that the pointing times for the unlimited roll rate model are always less than or equal to those for the nominal model with one exception. The heading angle/pitch angle pair of 179/45 is 0.01 s slower for the unlimited model. This is due to the computer tolerances and is practically insignificant.

Table 4.6. Pointing times for optimal trajectories of unlimited roll rate model

Heading Angle (deg)	Pitch Angle (deg)								
	-88	-60	-45	-20	0	20	45	60	88
15	3.69	2.53	1.88	0.88	0.56	0.96	2.07	2.77	4.33
45	3.71	2.88	2.44	2.00	1.88	2.08	2.67	3.19	4.36
90	3.78	3.78	3.81	3.89	4.00	4.13	4.30	4.38	4.46
135	3.84	4.79	5.32	6.17	6.79	7.06	6.50	5.84	4.56
165	3.87	5.24	6.06	7.66	9.05	9.84	8.11	6.67	4.59
179	3.87	5.29	6.25	7.96	9.59	10.88	8.36	6.77	4.60

4.3.3.2 Polar contour plots

Figure 4.27 demonstrates that the only advantage of the unlimited model comes in lower hemisphere pointing where the aircraft must roll more than 90 deg. The unlimited model is able to travel about 10 deg farther in the same time when the maneuver requires a 180 deg bank. This advantage occurs regardless of how long the trajectory lasts because the aircraft cannot begin its α increase until the bank is at least 90 deg. If it did not wait, the aircraft would pull into the upper hemisphere and away from its target.

4.3.3.3 Time histories

For lower hemisphere pointing, bank angles greater than 90 deg are required, and the unlimited roll rate model excels. This is because the aircraft cannot start using its angle-of-attack capability until its lift vector is properly oriented. It would be counterproductive to pull up immediately because the elevation angle is increasing instead of decreasing. So the aircraft waits until the roll is substantially completed. This is demonstrated in all the time histories of this section (Figures 4.28 through 4.34): angle of attack does not increase much until the bank angle is near the maximum.

Figures 4.28, 4.29, and 4.30 specifically show lower hemisphere pointing. Because the maximum bank angle is reached sooner, all the other action happens earlier, too. The time histories of the two models (except for roll rate) have generally the same shape for each target, but the unlimited model plots are compressed in time. In particular, g_{max} is reached sooner.

One notable exception to the same-shape time histories is that of thrust in Figure 4.30. The nominal model thrust decreases to about 35% of A/B and then increases to over

50% A/B. The unlimited roll rate model does not exhibit the increase. This is to keep the speed from increasing enough to violate the g_{max} limit. It is the same trade-off between angle of attack and speed we have seen earlier to optimize the load factor history. The trajectory is short enough that the high angles of attack used will not increase drag to the point that speed is lowered enough to degrade the load factor history.

In some of the runs, particularly to targets in the lower hemisphere, we see a strong oscillation in the bank angle history. This is similar to the Gibbs phenomenon seen in Fourier series [29]. It comes from the cubic spline fit of the control and state histories in the 9 segments of the trajectory. Because the model has infinite roll rate capability, the optimization is trying to bank to the necessary angle in zero time. This cannot be accomplished because of the finite length segments. The bank angle therefore increases as fast as it can with the cubic spline fit. In Figure 4.31, we see what happens when we increase the number of segments to 19 in a trajectory to (45,45). The bank angle history looks a bit more like a step, the oscillations decrease, and the final time is lowered from 2.67 s to 2.47 s.

One of the goals of the research was to see how much roll rate the model would use given an infinite capability. The two highest roll rates ever used were 739 deg/s for (15,0) and 584 deg/s for (15,-45). The next highest rate was about 300 deg/s, exhibited in many runs. This is almost twice the nominal limit of 180 deg/s. Because of the cubic spline fit, we assume that the optimization is actually using less roll rate than it wants. Indeed the maximum roll rate for (45,45) with 19 segments increased from 68.2 deg/s to 73.3 deg/s. We would expect that in the limit as the segments became infinitely small, the roll rate would approach infinity. (By the same token we would also expect that the spool time would approach zero for the unlimited spool time cases as the trajectory was divided finer. Pitch rate would not necessarily approach infinity because the resulting high angles of attack carry drag and speed penalties)

Figures 4.32, 4.33, and 4.34 show cases in which infinite roll rate capability is not of significant advantage. In fact, the histories displayed in Figure 4.32 are almost identical. This is a short trajectory with little bank angle required (the target is (15,45)); so one would expect the histories to be similar. The very short time to point is expected because about half of the trajectory is spent at g_{max}. The thrust history was explained in Section 4.4.3.1.

(165,45) in Figure 4.33 also requires comparatively small bank angles, less than 90 deg; so the histories should be similar. The bank angle is actually less for the unlimited roll rate model for most of the trajectory. The purpose of a bank is to reorient the lift vector, so high bank angles increase the horizontal lift component. This requires a higher angle of attack to maintain the vertical lift component needed to achieve the desired elevation angle and to balance the weight. The higher angle of attack carries with it drag, speed, and load

factor penalties. The longer the aircraft remains at high-α conditions, the more the load factor is adversely affected. The bank angle for the unlimited model, therefore, nearly steps at the end to acquire the desired target without the necessity of high-α for a long time.

The bank angle required for (90,0) is about 90 deg, so the turn is nearly horizontal. This suggests that the roll rate limits are not too important, and indeed Figure 4.34 shows this. Aside from the bank angle, the chief difference in the histories is in the thrust plot. The unlimited model thrust increases almost immediately, while the limited model thrust remains nearly constant. This is due to the tradeoffs in angle of attack and speed to keep the aircraft at g_{max}. One would not want too much angle of attack, or the nose would pull up from a horizontal turn and away from the desired elevation. To allow a lower α, speed must be increased to keep the lift and the load factor as high as possible. This is done by throttling towards A/B. If the limited model thrust did the same, the lift would increase enough to violate the load factor constraints. This run also exhibits the bank angle oscillations described earlier.

The results for this model suggest that the simple metric of time to roll 180 deg may be a valid measure of aircraft agility.

4.4 Agility Comparison Using Standard Maneuvers

4.4.1 Time matrix

Times for the unlimited pitch rate model performing standard maneuvers are shown in Table 4.7. These times are all less than those for the nominal model performing standard maneuvers (Table 4.2). With one exception the times are all greater than or equal to the times for the unlimited pitch rate model performing optimal maneuvers (Table 4.4). (The exception is for (15,60) where there is a 0.01 s difference attributable to the finite tolerances of the digital computer.) Thus we would expect that using standard maneuvers to compare agility actually disguises the true capability of the aircraft in some regions.

4.4.2 Polar plot

Two polar contour plots are given in Figure 4.35. The first plot (Figure 4.35(a)) shows the difference between the nominal and unlimited pitch rate models in standard maneuvers. The second polar plot (Figure 4.35(b)) compares the optimal and standard maneuvers of the unlimited pitch rate model.

We can see that the standard maneuvers disguise much of the extra agility of the unlimited pitch rate model. According to the 8 s contour of Figure 4.35(a), in fact,

Table 4.7. Pointing times for standard trajectories of unlimited pitch rate model

Heading Angle (deg)	Pitch Angle (deg)								
	-88	-60	-45	-20	0	20	45	60	88
15	4.98	4.14	2.89	1.53	0.73	0.94	2.09	2.18	3.59
45	4.92	4.21	3.82	3.05	2.52	2.14	2.64	2.90	3.64
90	4.87	4.27	3.87	3.67	3.77	3.96	4.21	4.34	3.76
135	4.88	4.98	5.55	6.71	7.48	7.93	7.49	6.66	3.89
165	4.93	5.59	6.83	8.61	9.82	10.86	8.29	6.78	3.93
179	4.96	5.77	7.37	9.40	10.74	11.94	8.29	6.71	3.94

there is very little difference between the unlimited and nominal models. But Figure 4.35(b) shows we can achieve an extra 10 deg in the same time merely by switching to the optimal trajectory. In the standard maneuver, the airplane is constrained to roll before pitching regardless of the pitch rate limits. This takes away much of the advantages. The only time the standard maneuver accurately displays agility is in the split-S where no roll is required.

There is a curious effect in the 4 s contour of the unlimited standard maneuver. Figure 4.35(b) shows that the contour approaches the unlimited optimal contour as the maneuver plane approaches 180 deg, and Figure 4.35(a) shows that it quickly switches to approach the nominal standard contour. The region in question surrounds the point directly below the aircraft at an elevation angle of -90 deg. This point necessitates a full 180 deg bank, and therefore the standard maneuver requires the aircraft to wait a long time before pitching. The optimal trajectory places no such constraints on the model, so it is free to pitch and roll simultaneously. This effect is not noticeable in the other contours because roll angles much less than 180 deg are required.

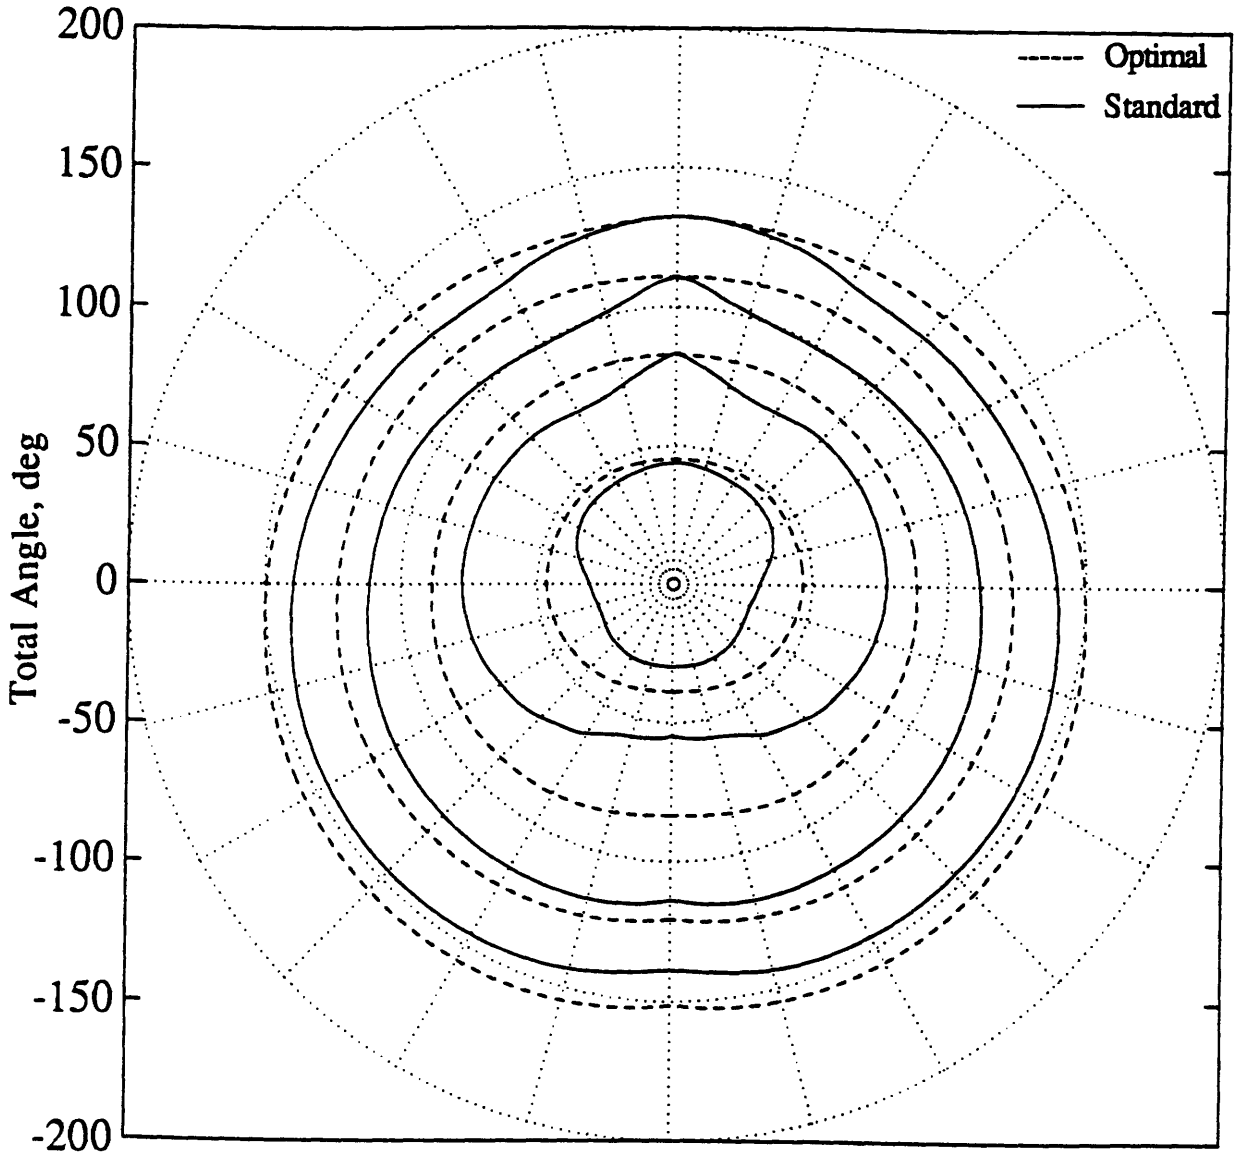

Figure 4.2. Contour plot of standard vs. optimal trajectories of nominal model. The only region with no difference in pointing time is near the 0 deg maneuver plane where no bank is required. The switch to the split-S is clear in the standard contours.

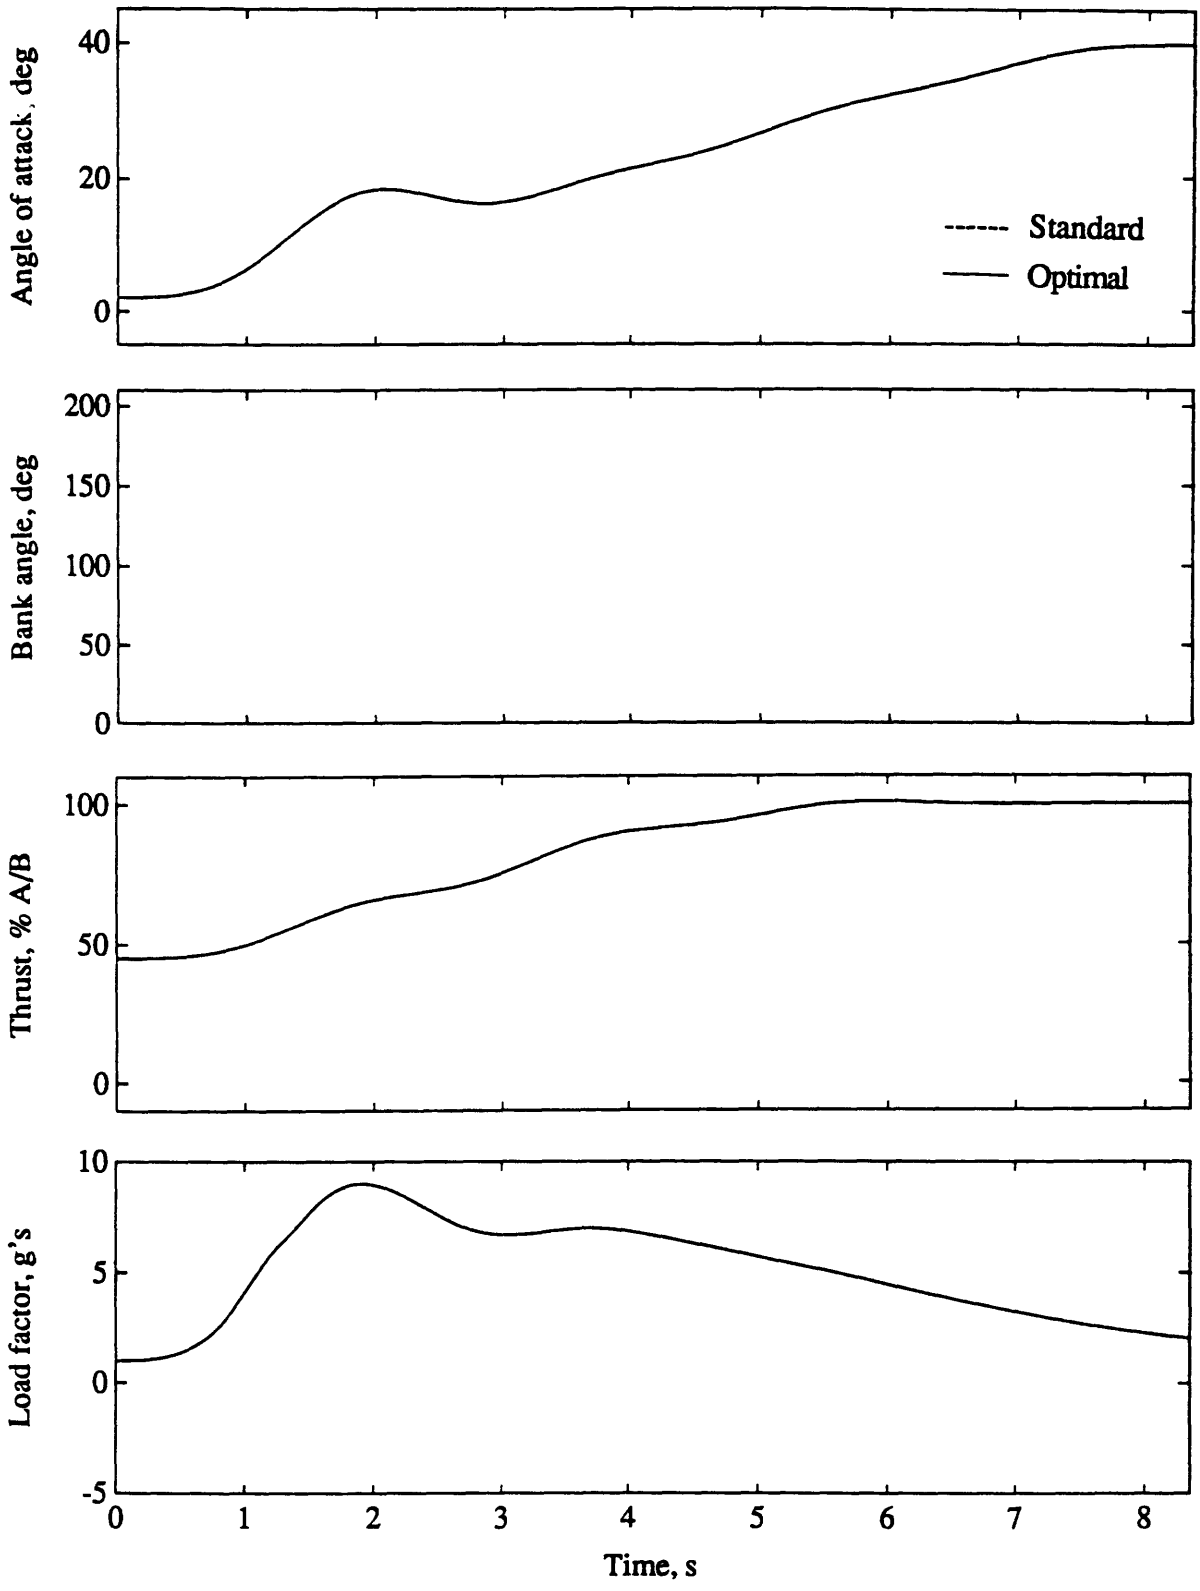

Figure 4.3. Standard vs. optimal trajectories for nominal model to (179,45). They are identical because no bank is required before the pull-up.

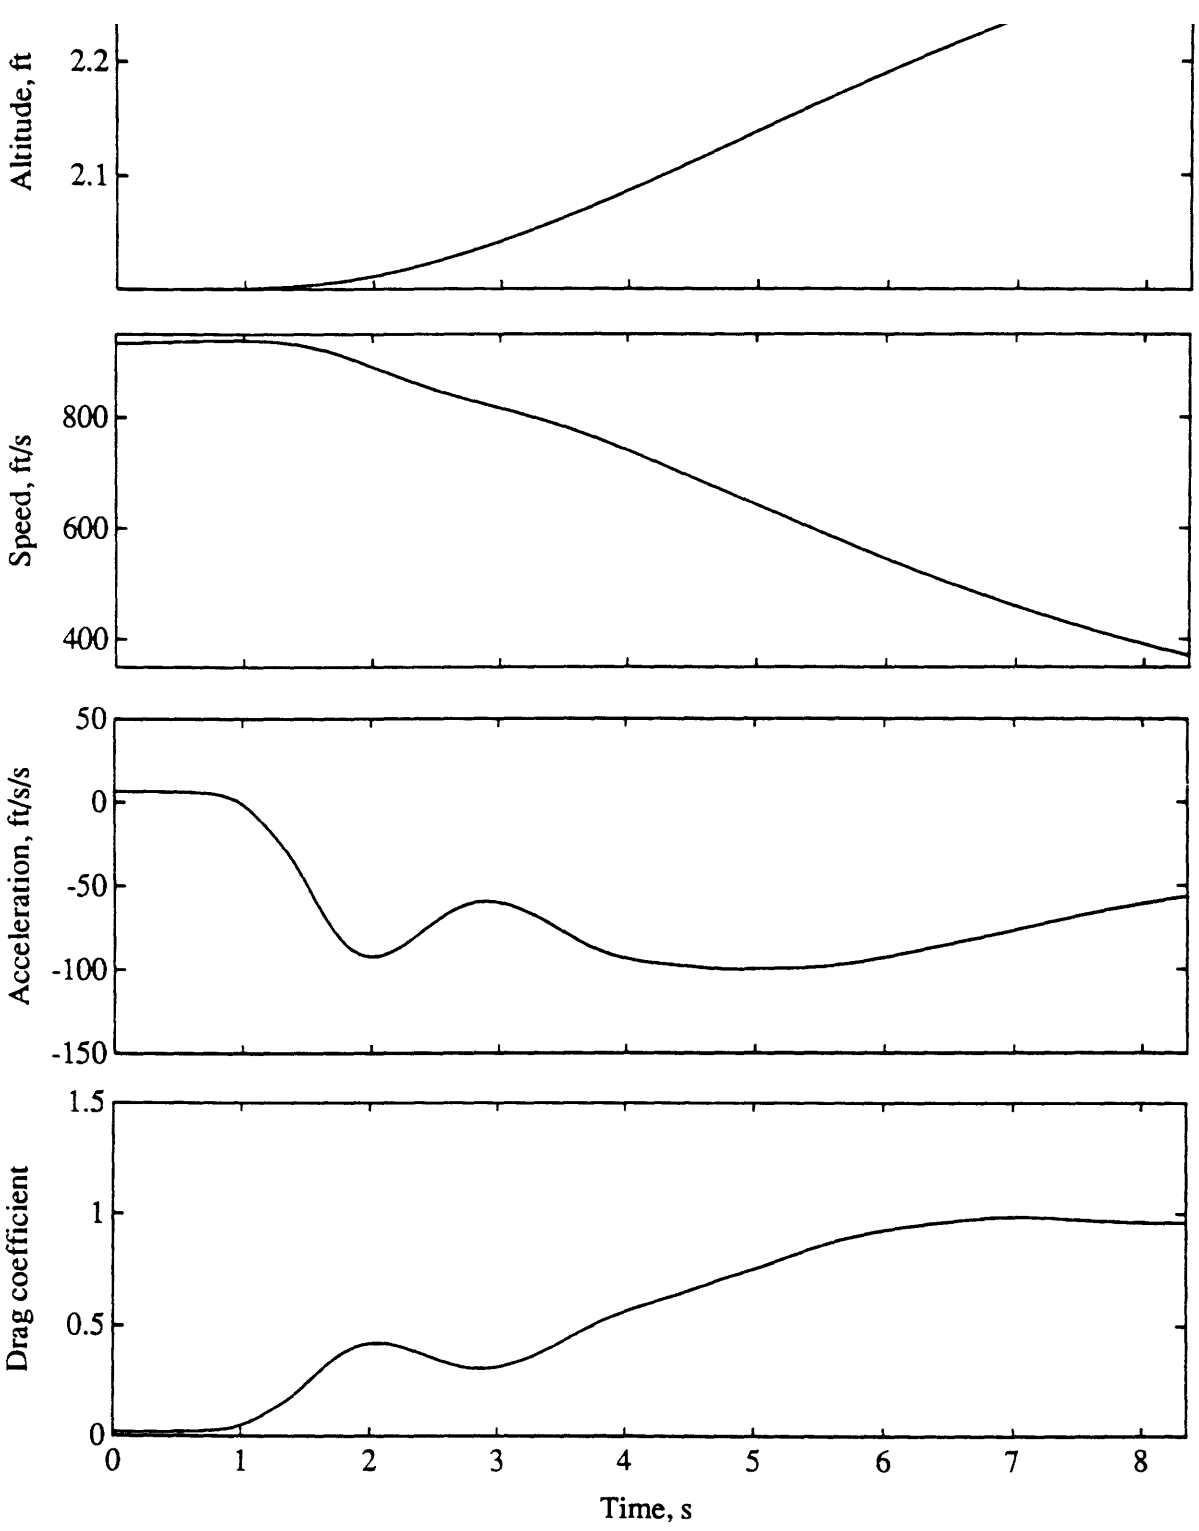

Figure 4.3 (cont'd)

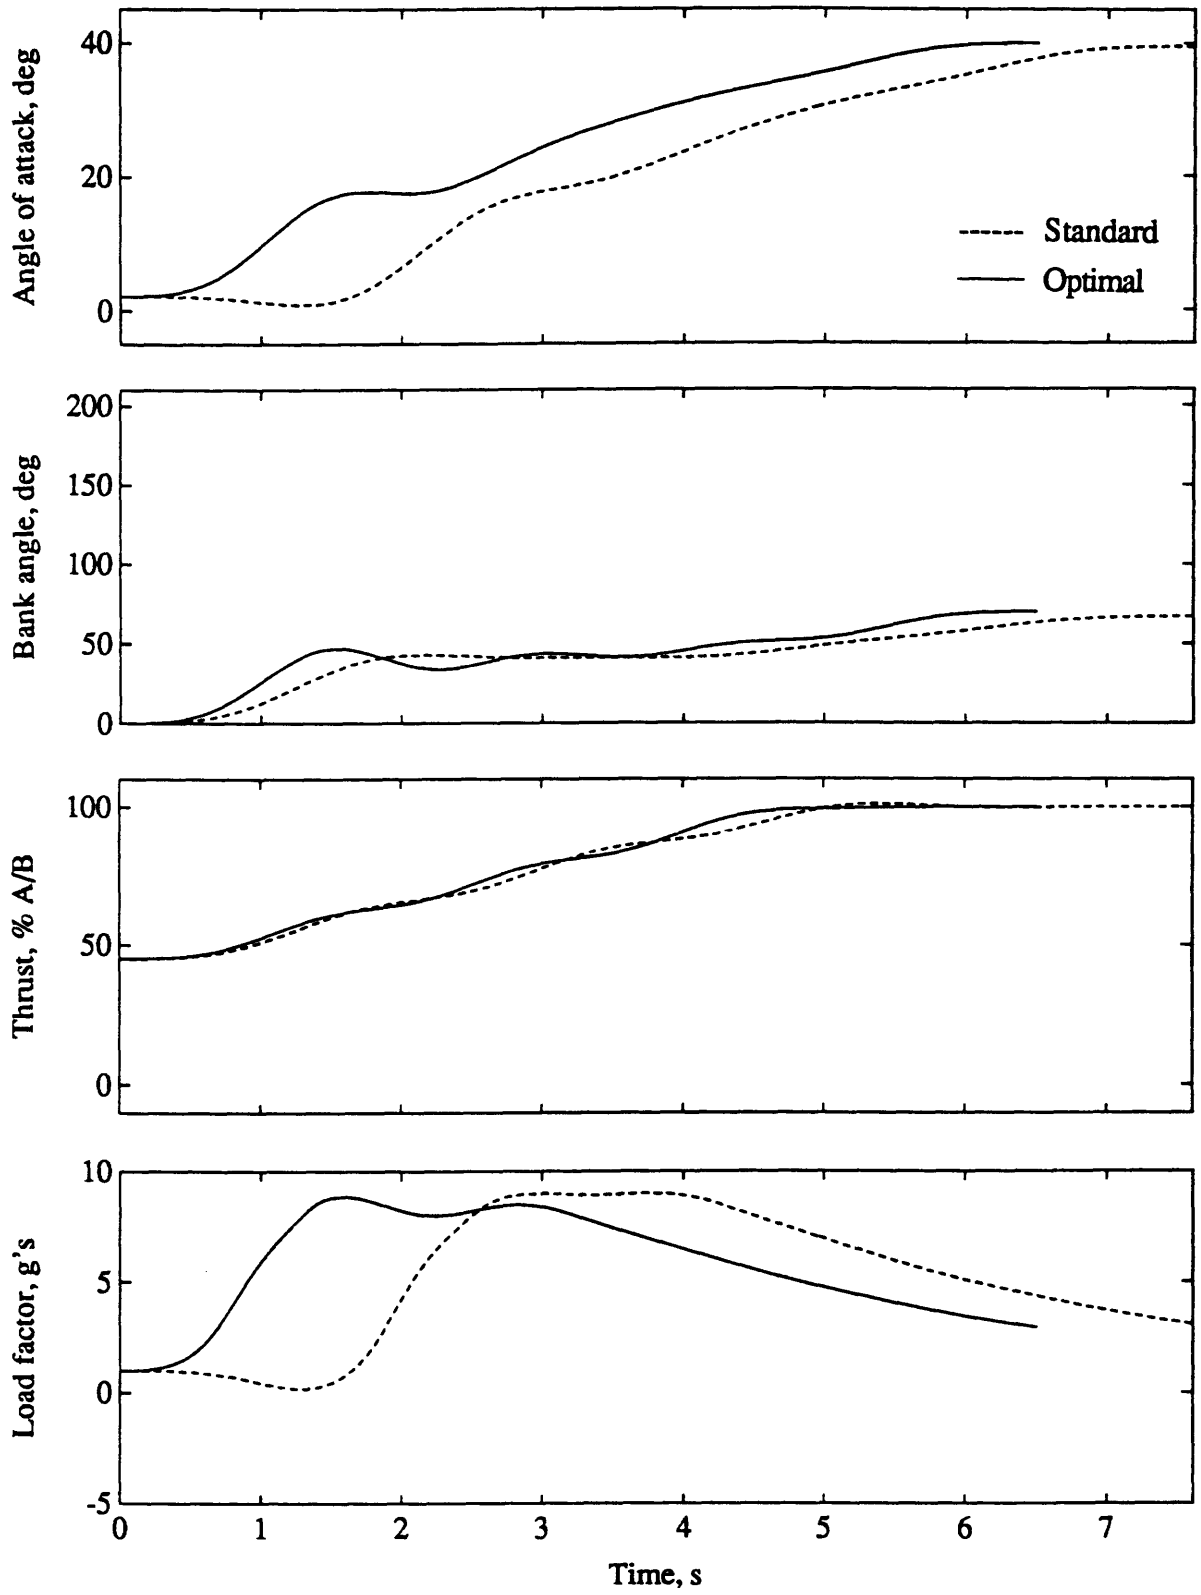

Figure 4.4. Standard vs. optimal trajectories for nominal model to (135,45). The standard maneuver is slower because the aircraft must wait for the roll to complete before it can begin the pull-up.

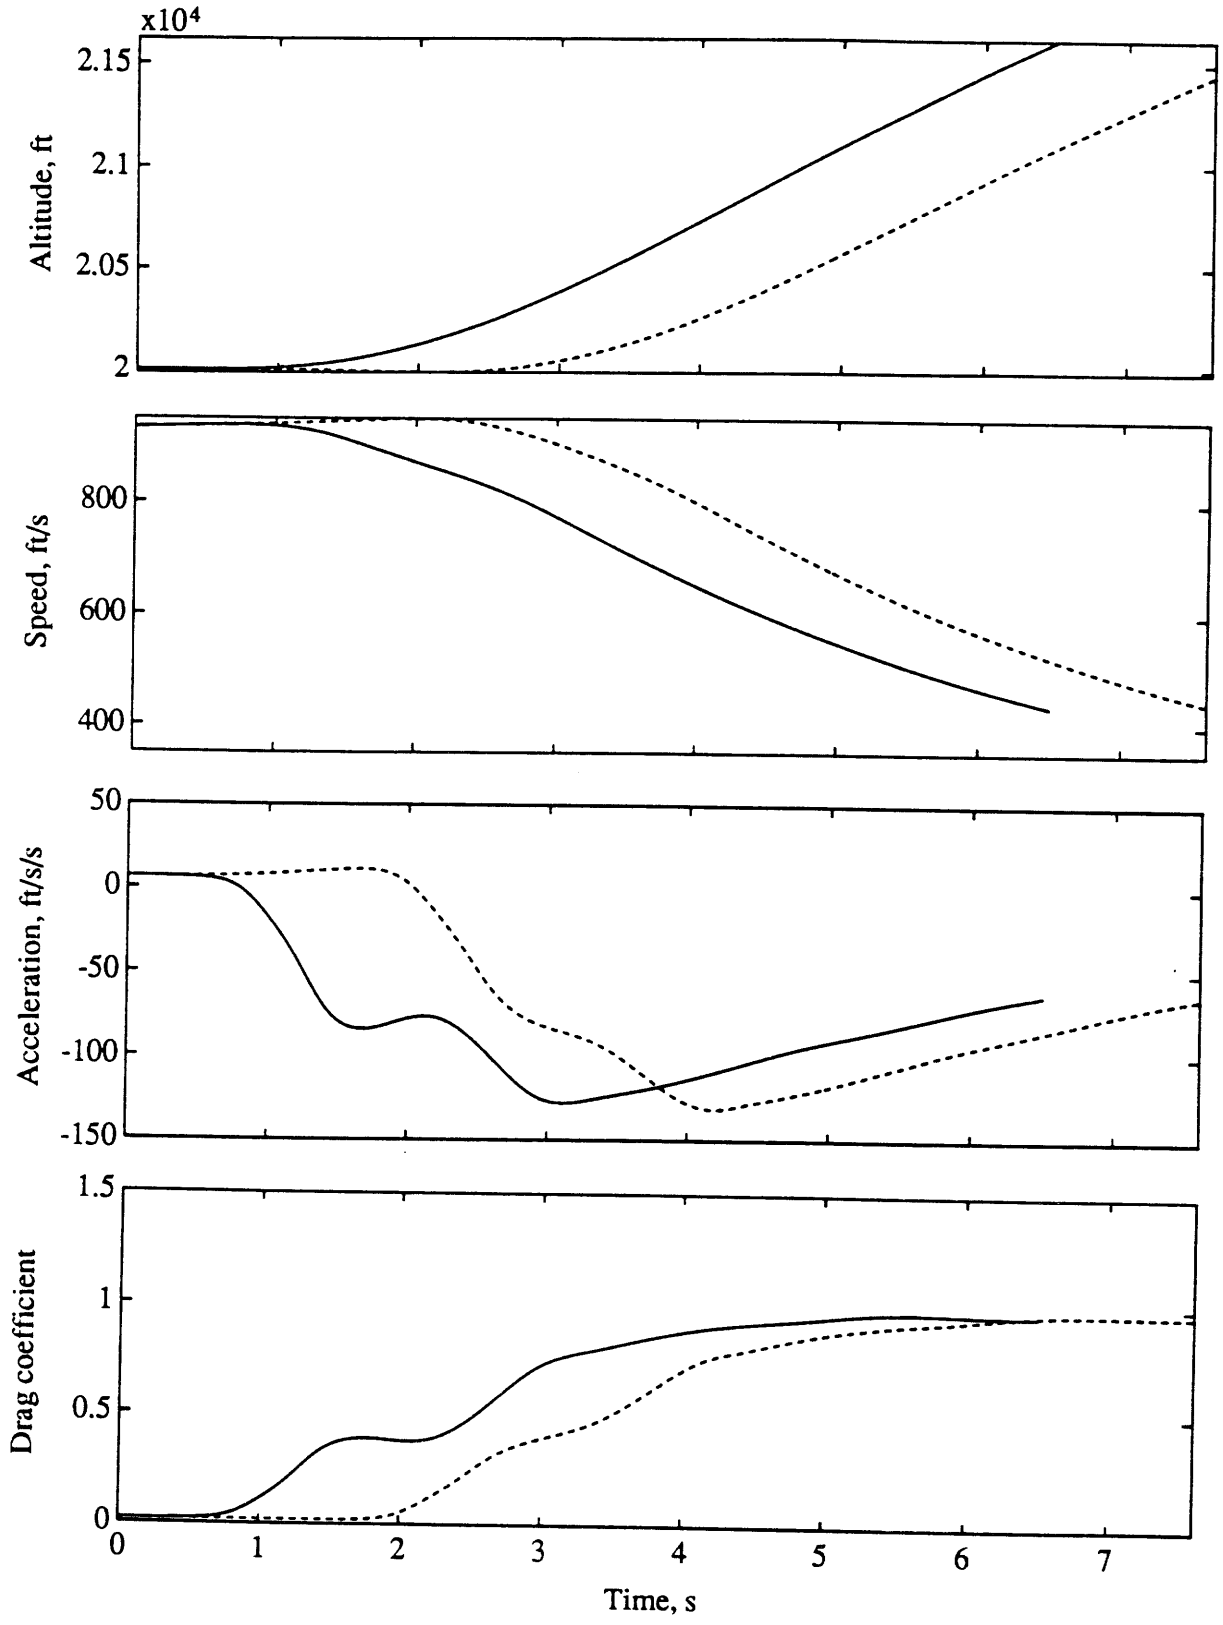

Figure 4.4 (cont'd)

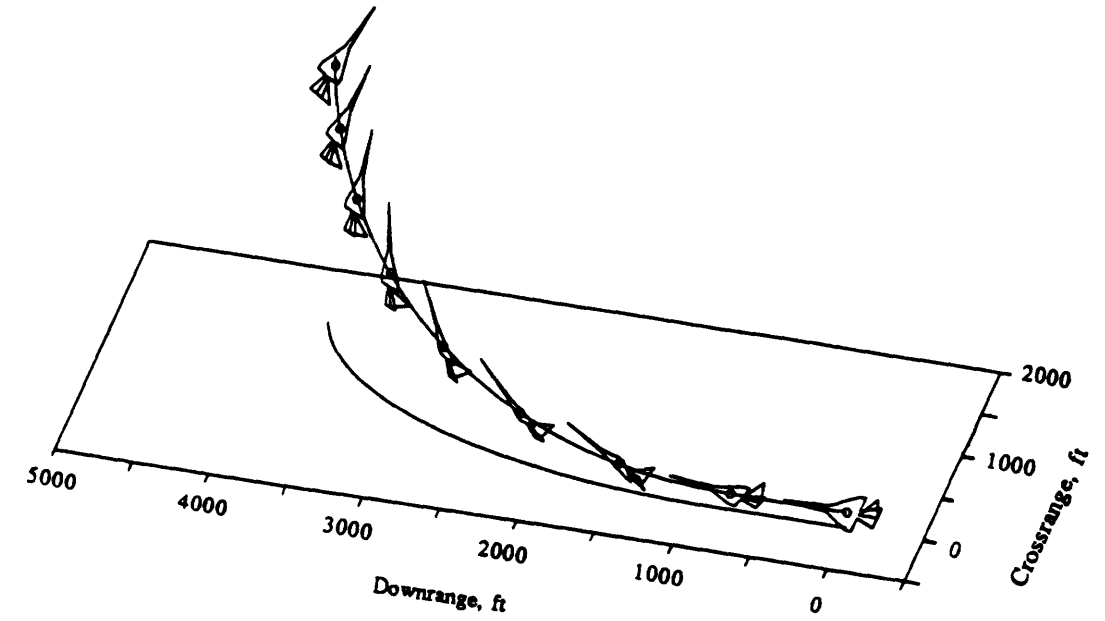

(a)

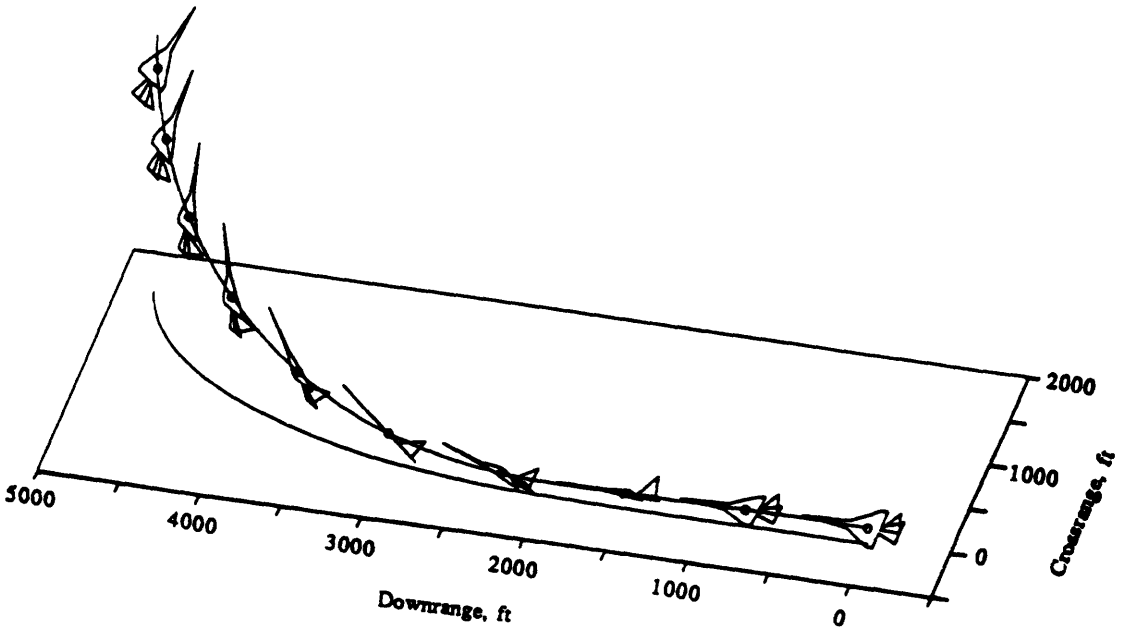

(b)

Figure 4.5. Three-dimensional trajectories to (135,45). The airplanes are drawn every 0.8 s. The optimal trajectory in (a) shows a simultaneous bank and pitch The standard maneuver in (b) waits until the bank is essentially complete before the pitch begins. Optimal downrange is much shorter than standard, but crossranges and final altitudes are nearly the same.

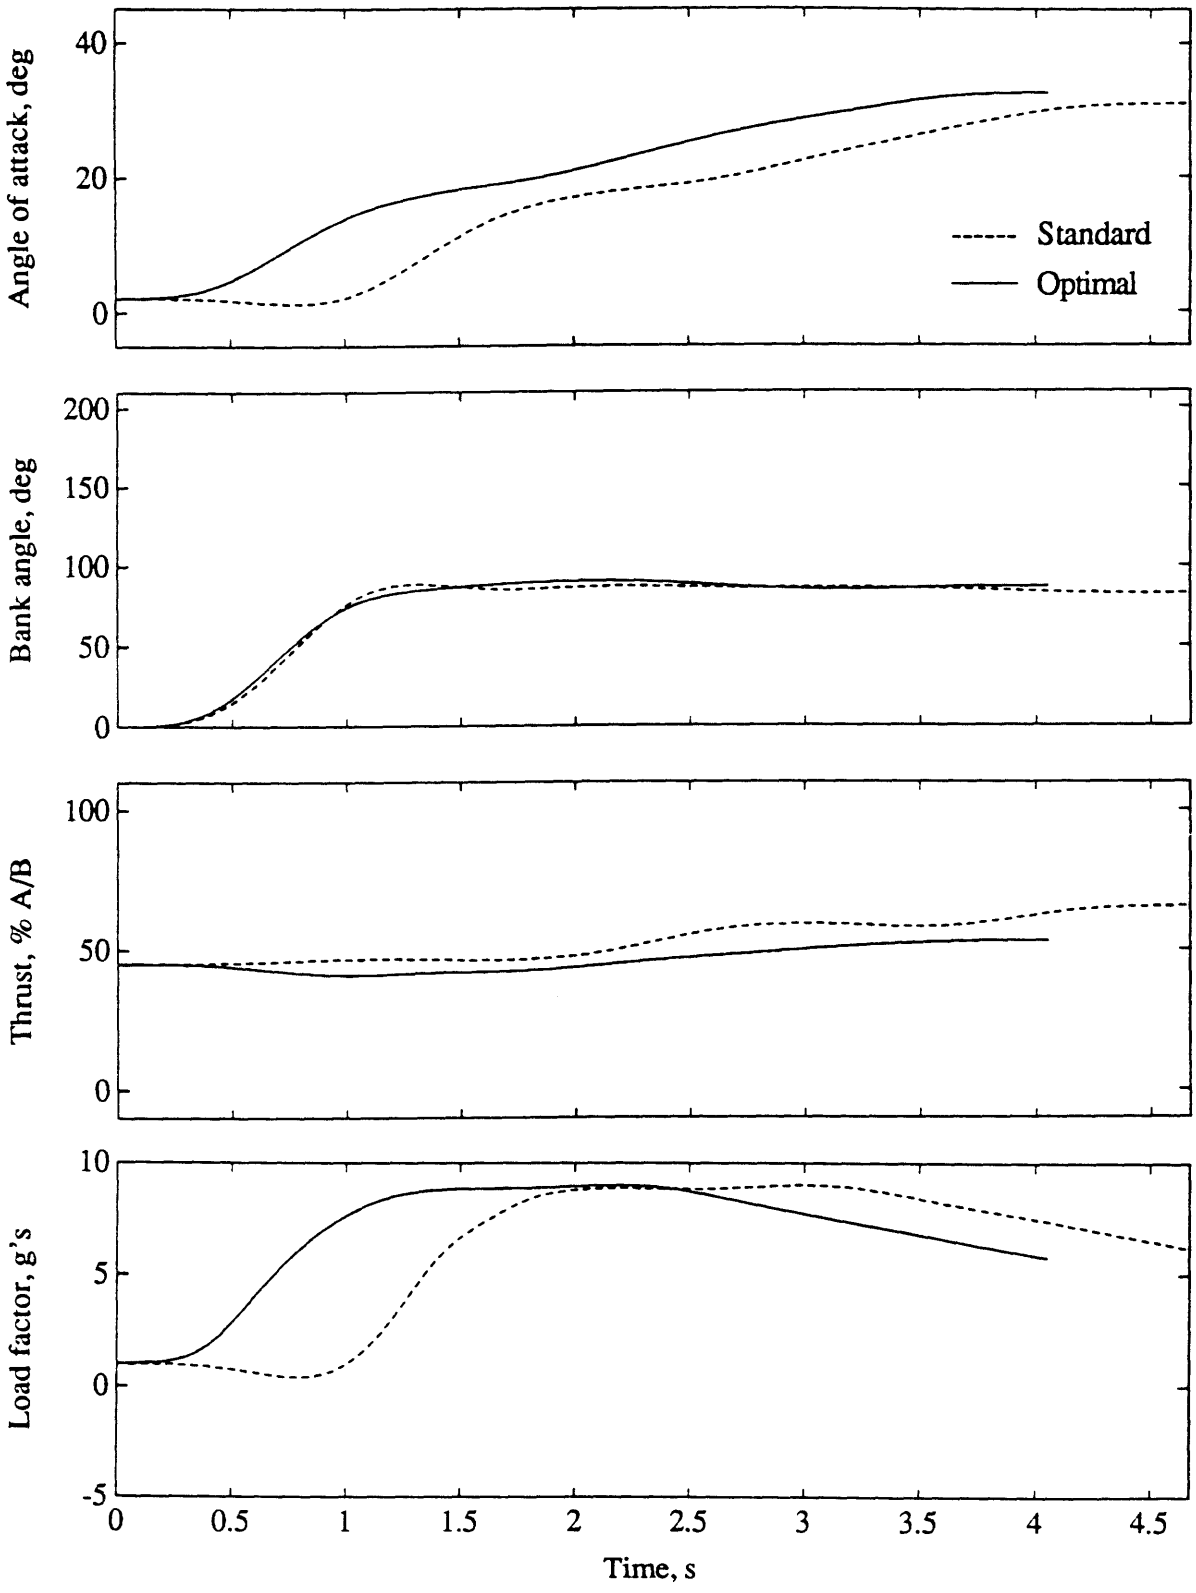

Figure 4.6. Standard vs. optimal trajectories for nominal model to (90,0). Thrust does not reach A/B in this run, in contrast to the previous runs. Maximum angle of attack is not reached because of pitch rate limits.

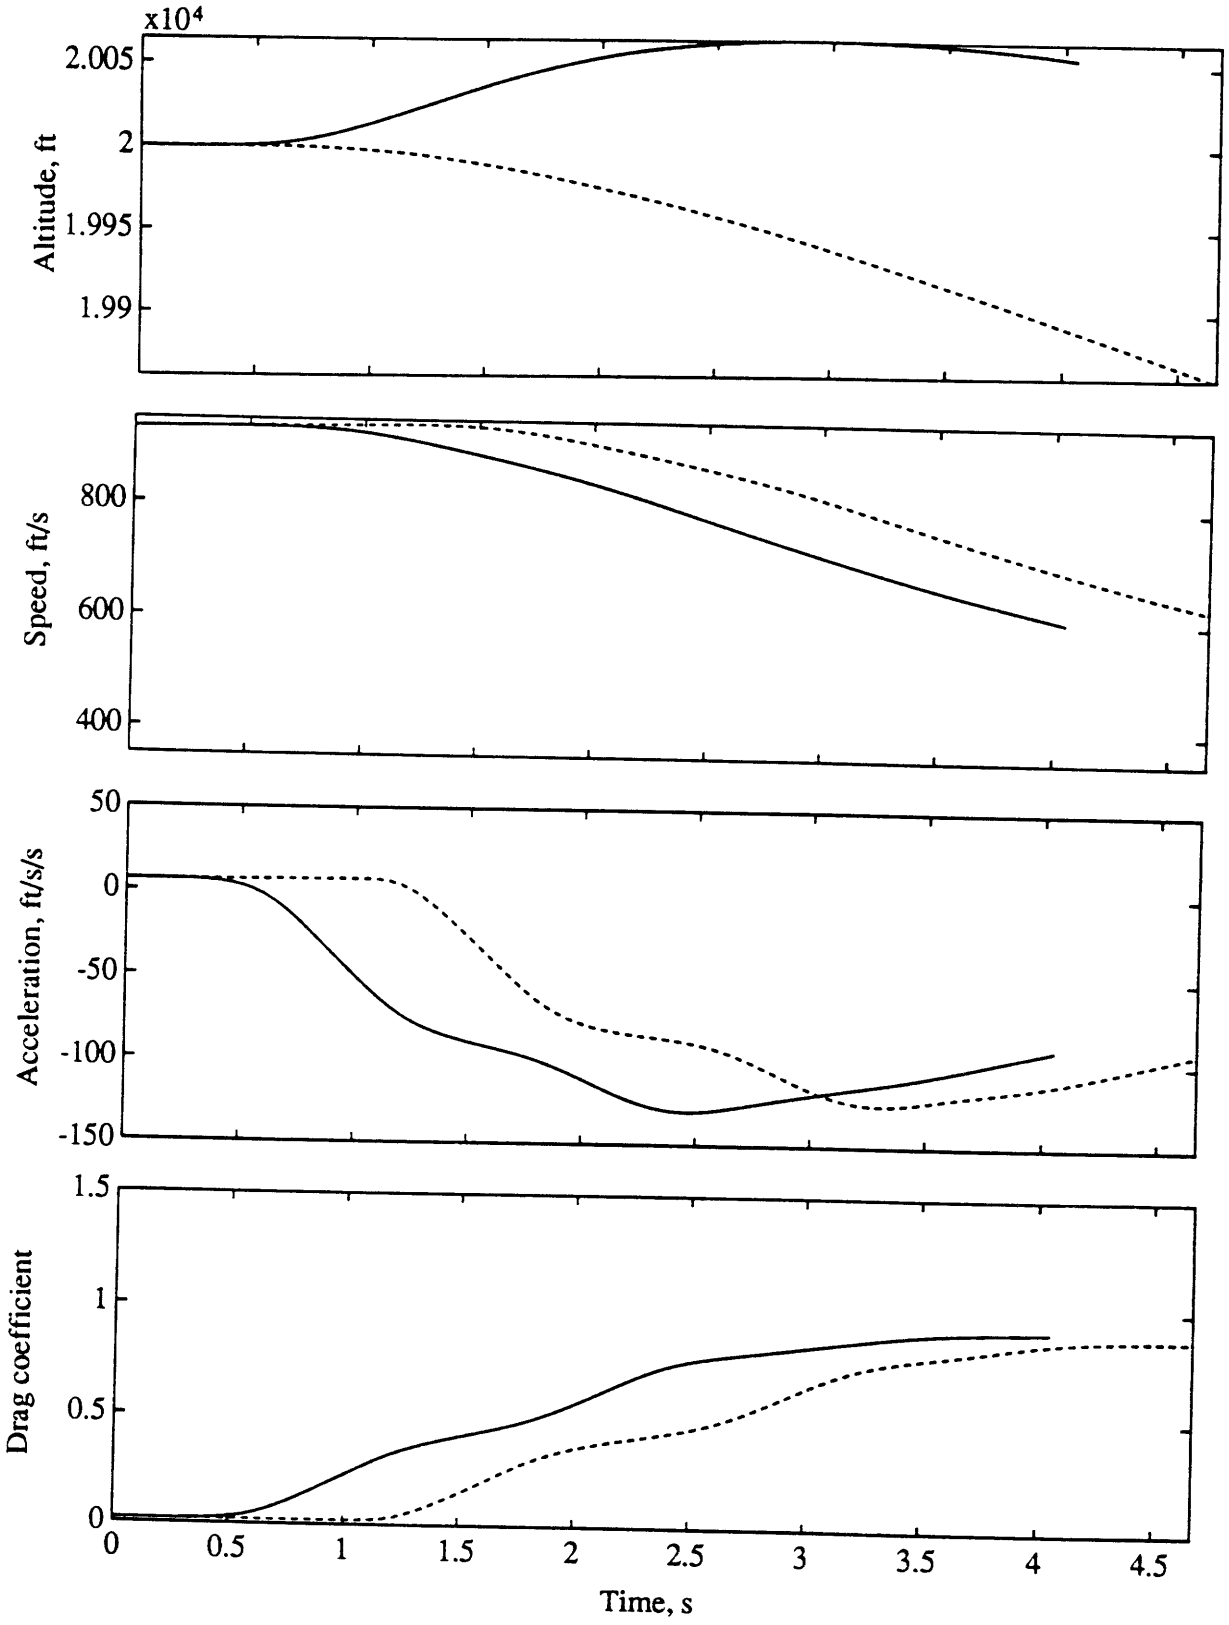

Figure 4.6 (cont'd)

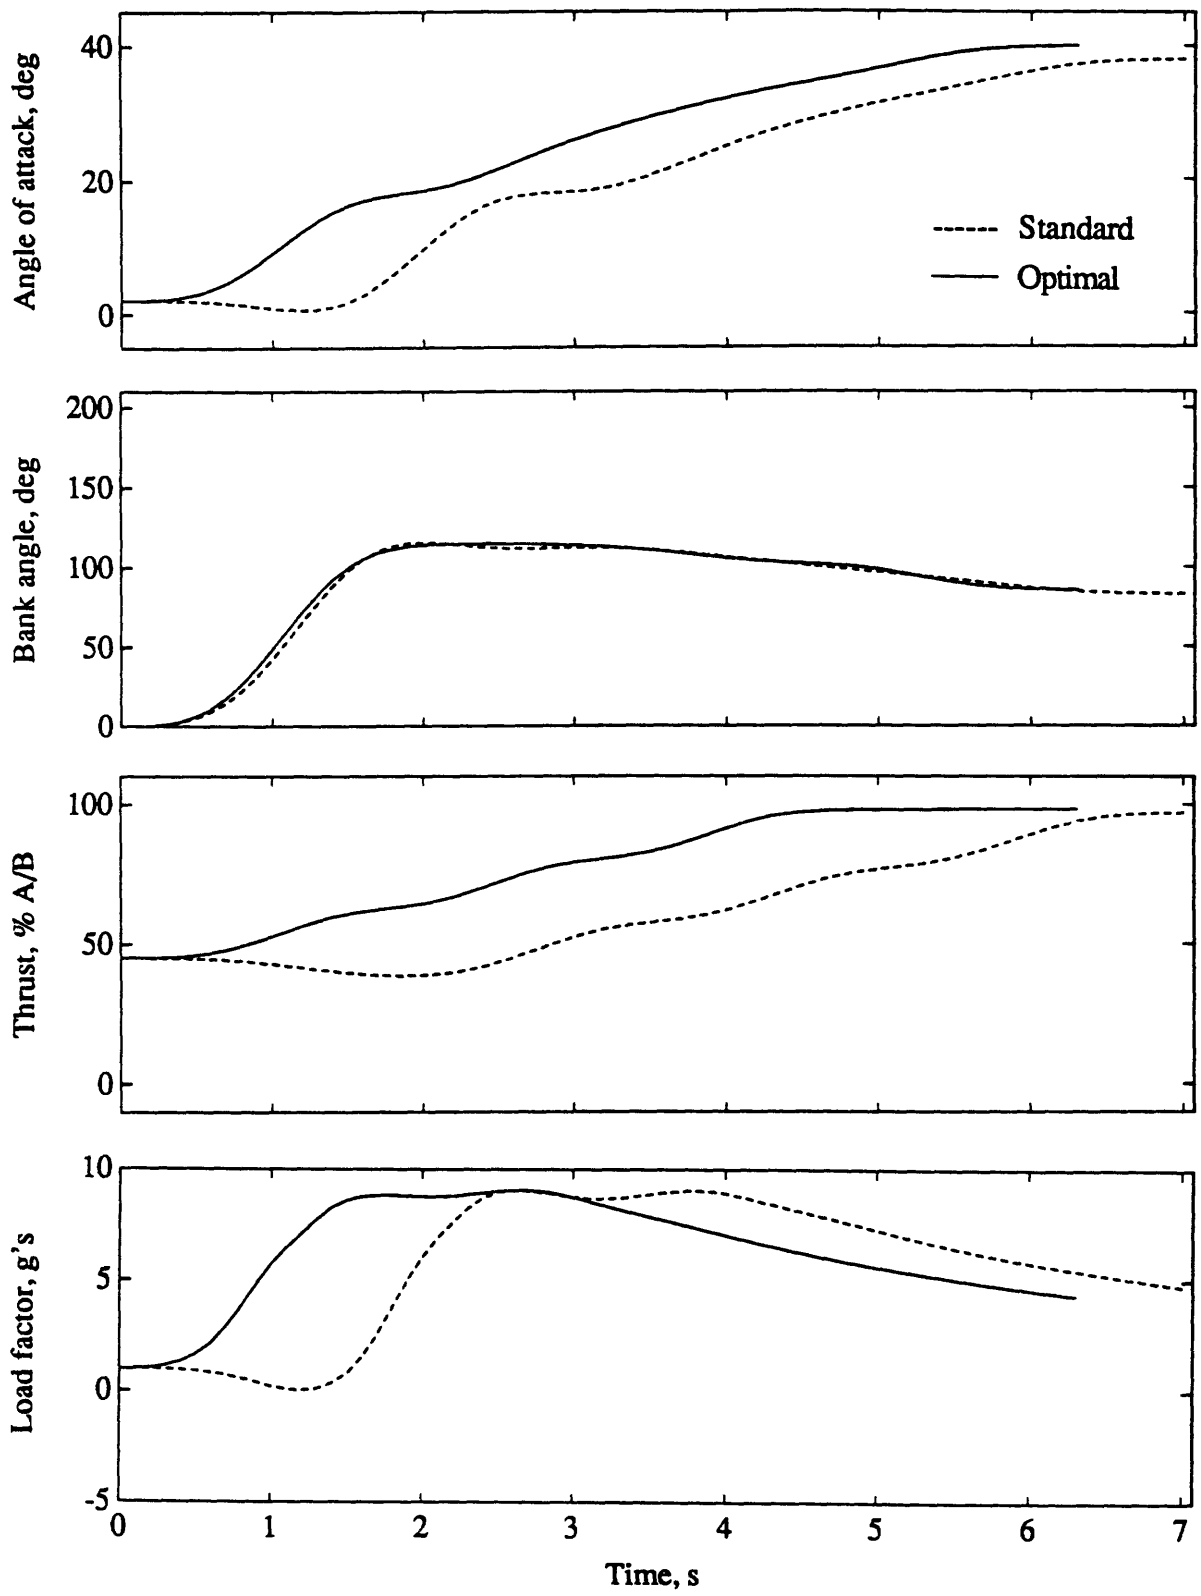

Figure 4.7. Standard vs. optimal trajectories for nominal model to (135,-20). Optimal thrust is much higher than standard because the higher α causes speed to bleed faster.

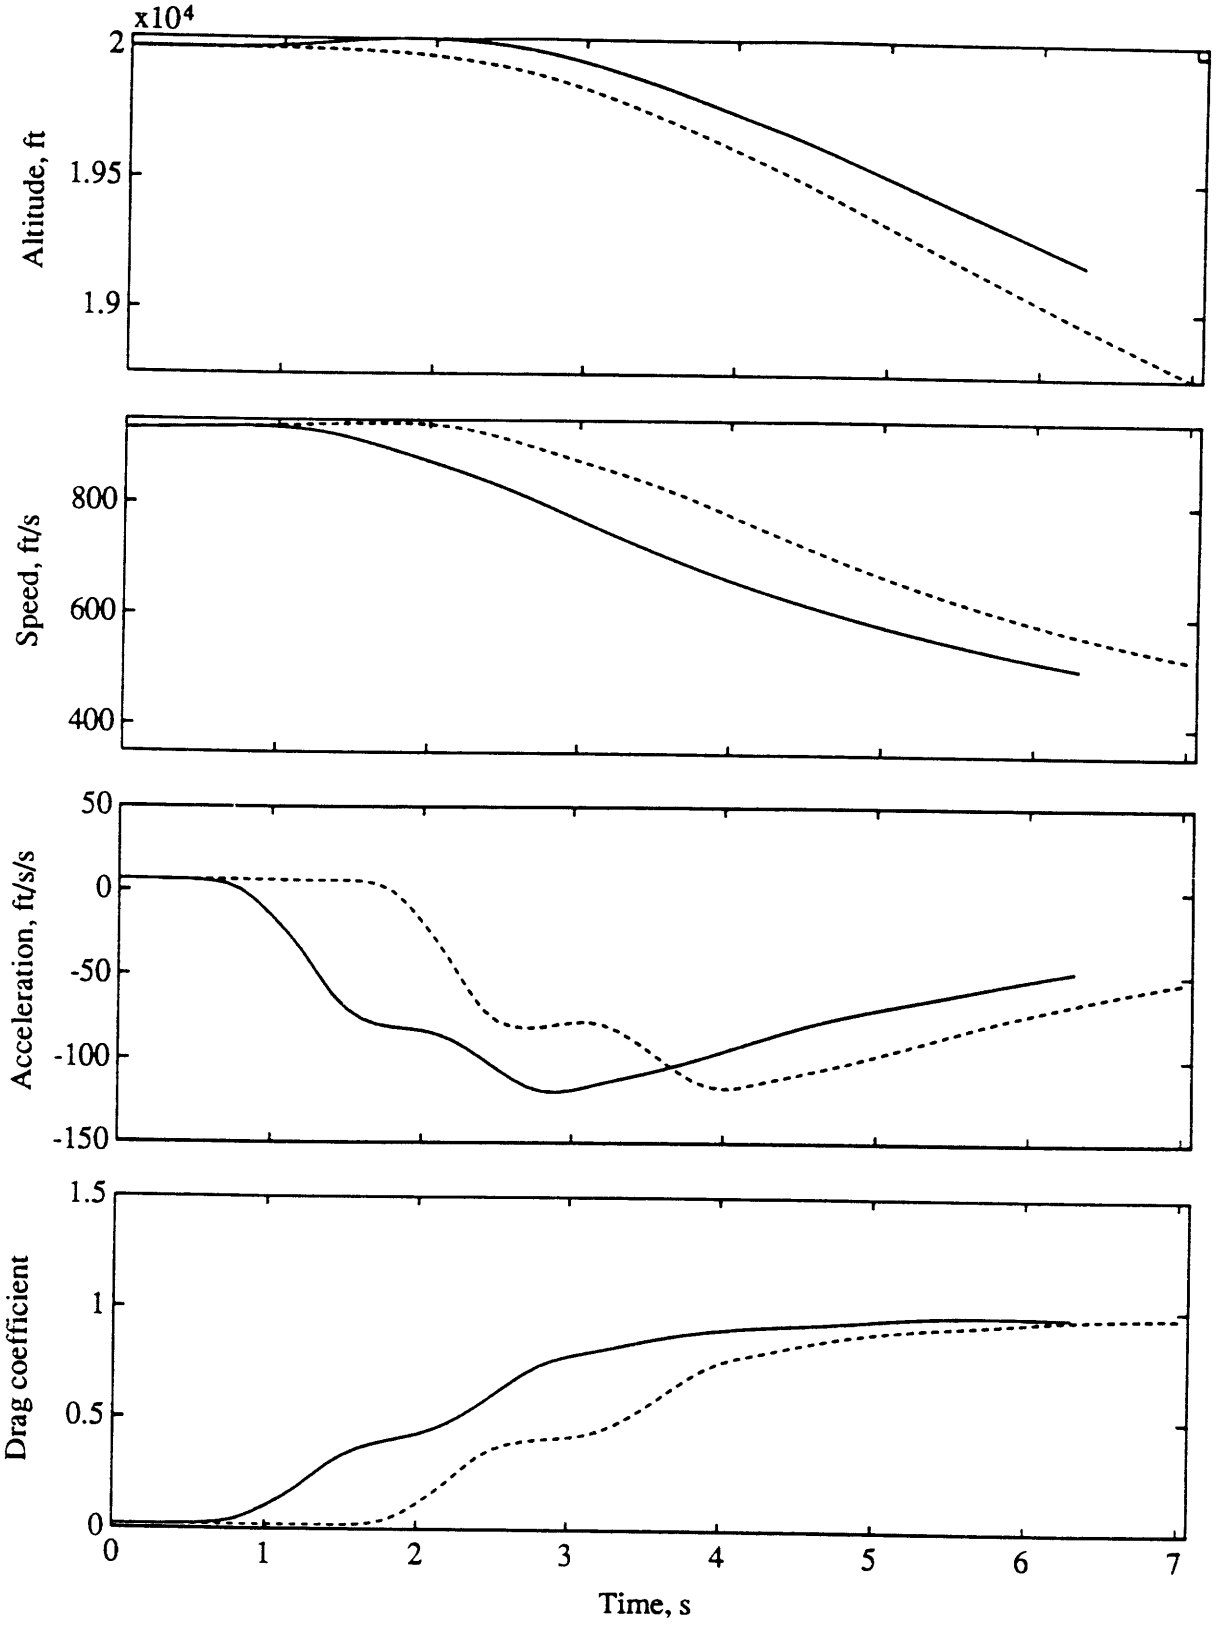

Figure 4.7 (cont'd)

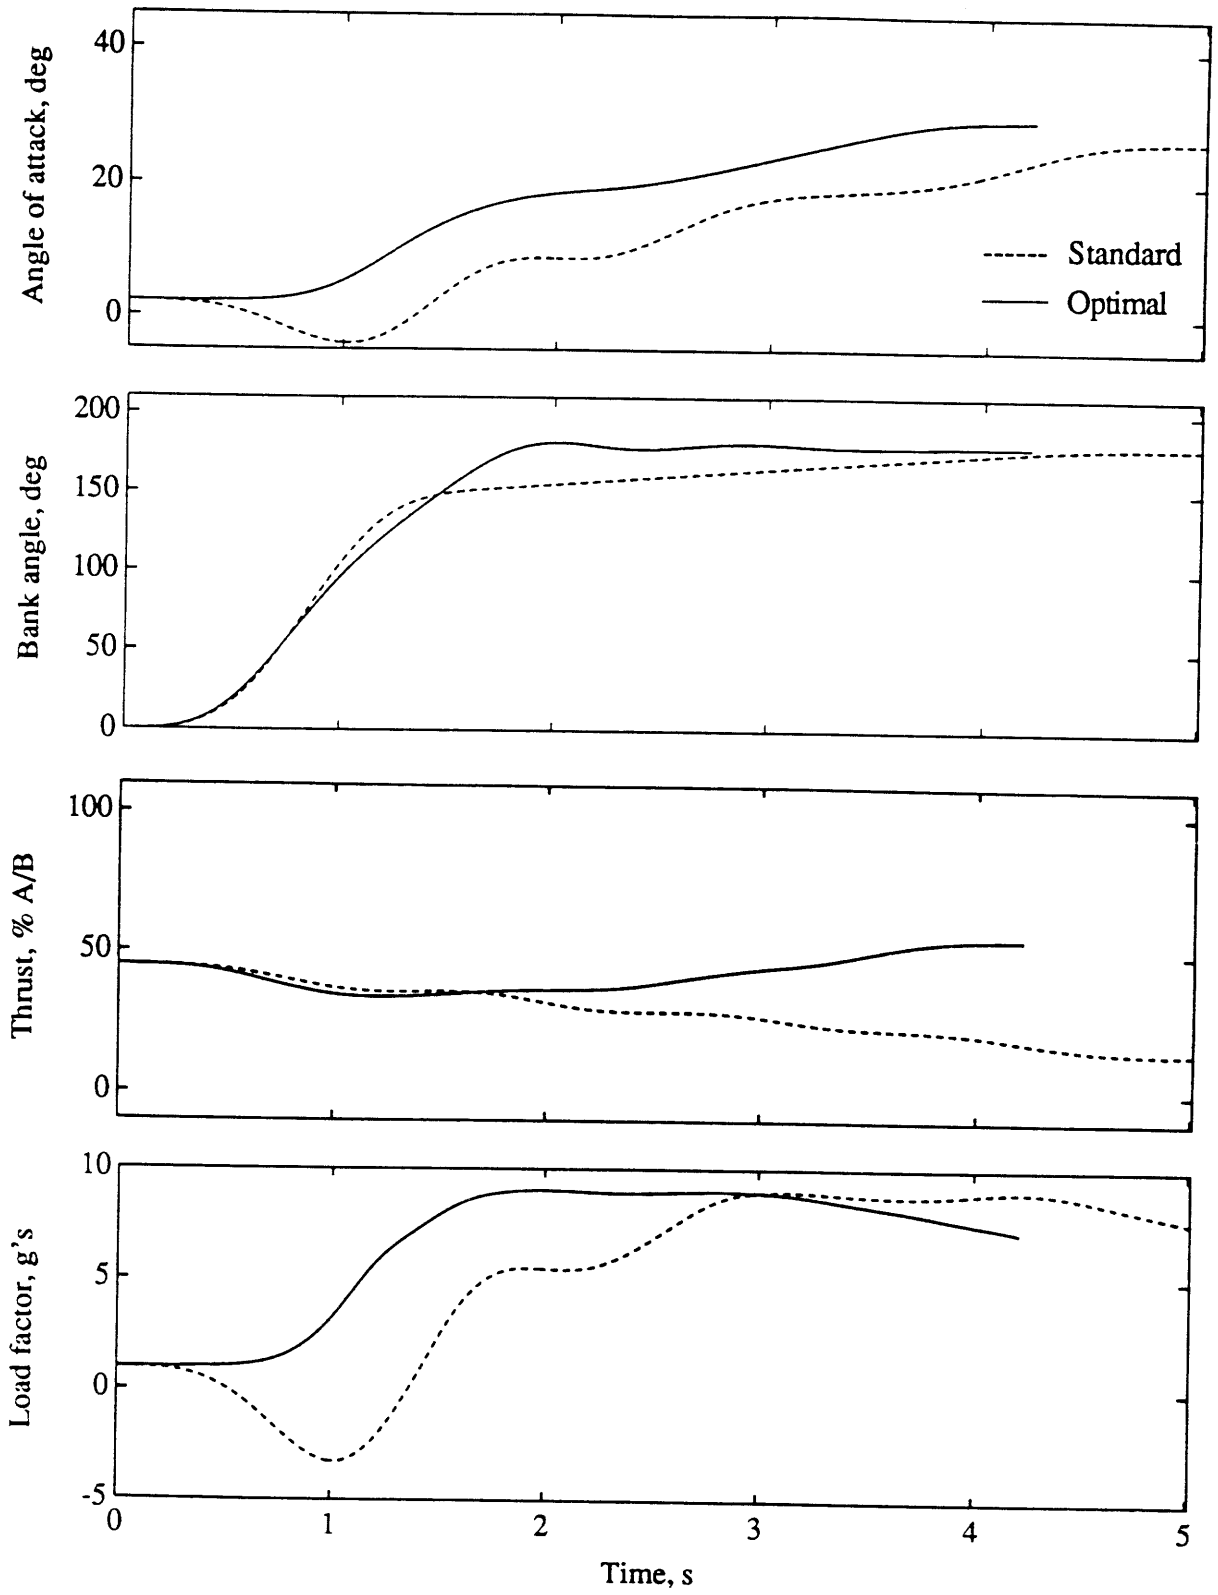

Figure 4.8. Standard vs. optimal trajectories for nominal model to (15,-88). Because the target is directly beneath the airplane and rolling is permitted only below 1 g for the standard maneuver, the standard α goes negative to get a headstart.

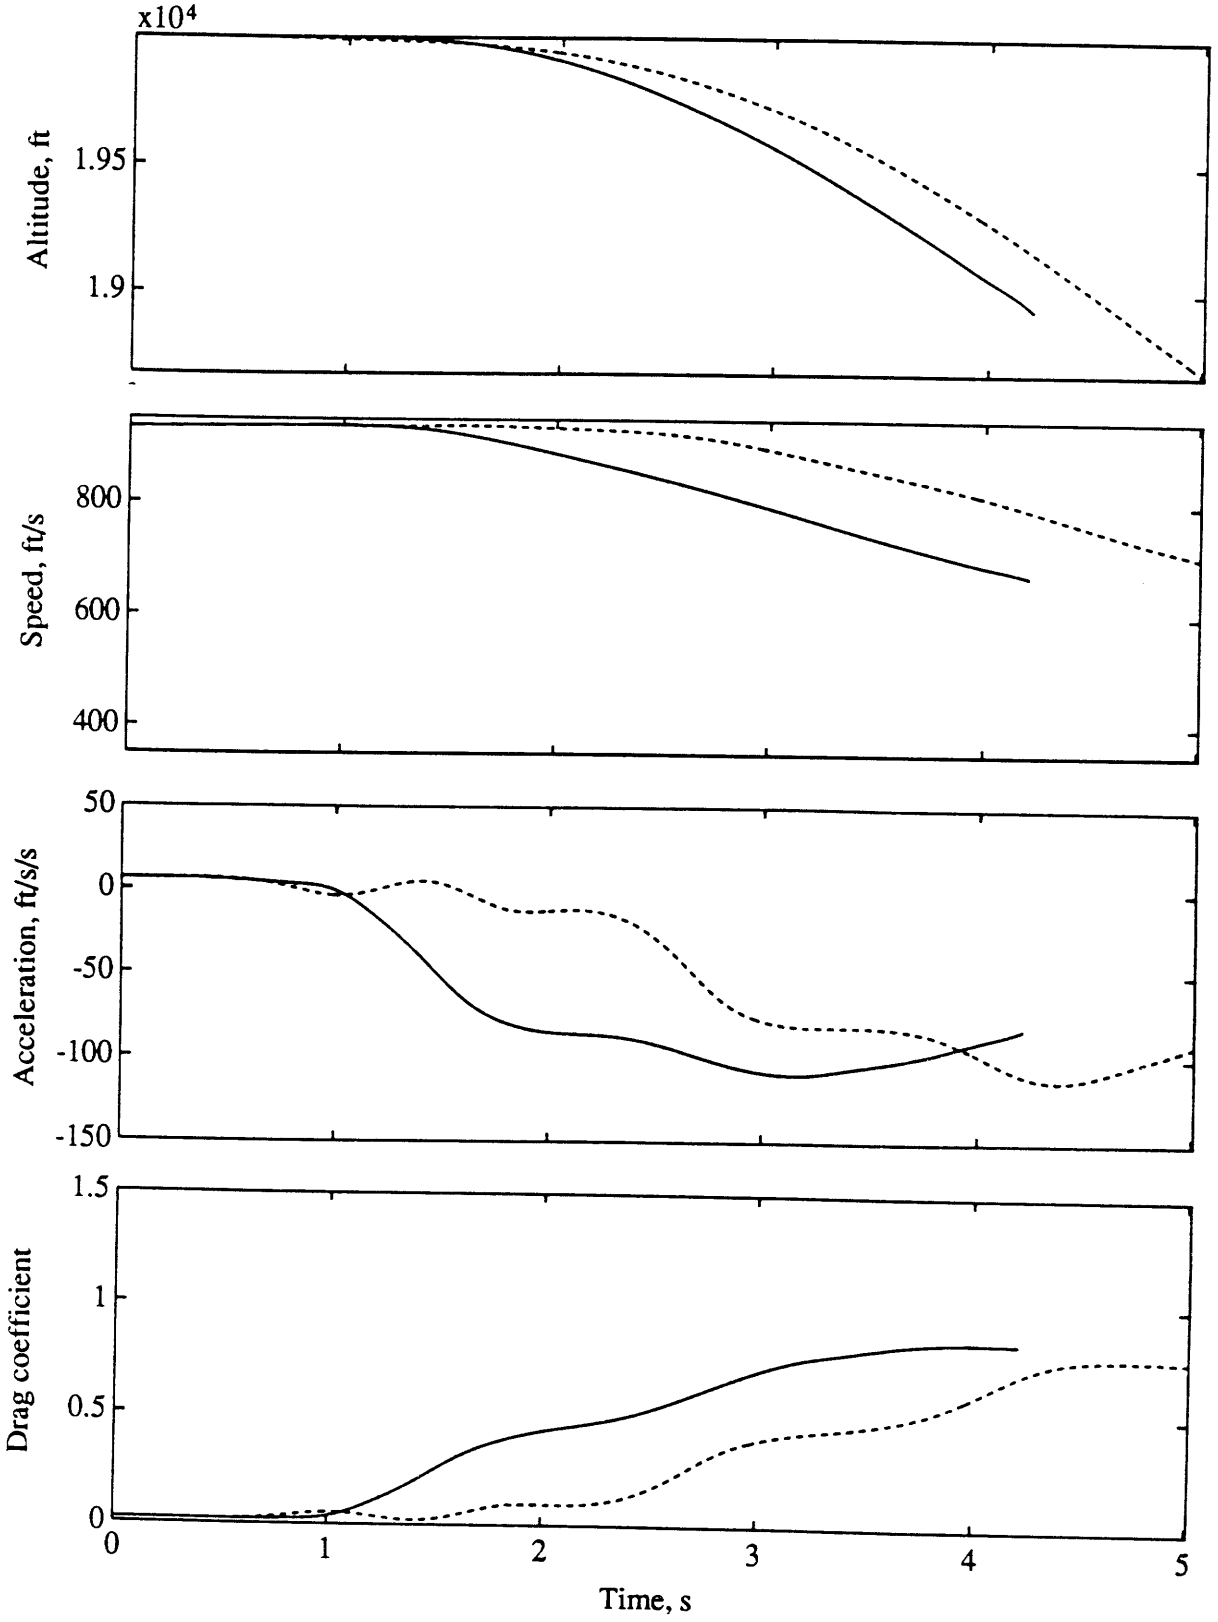

Figure 4.8 (cont'd)

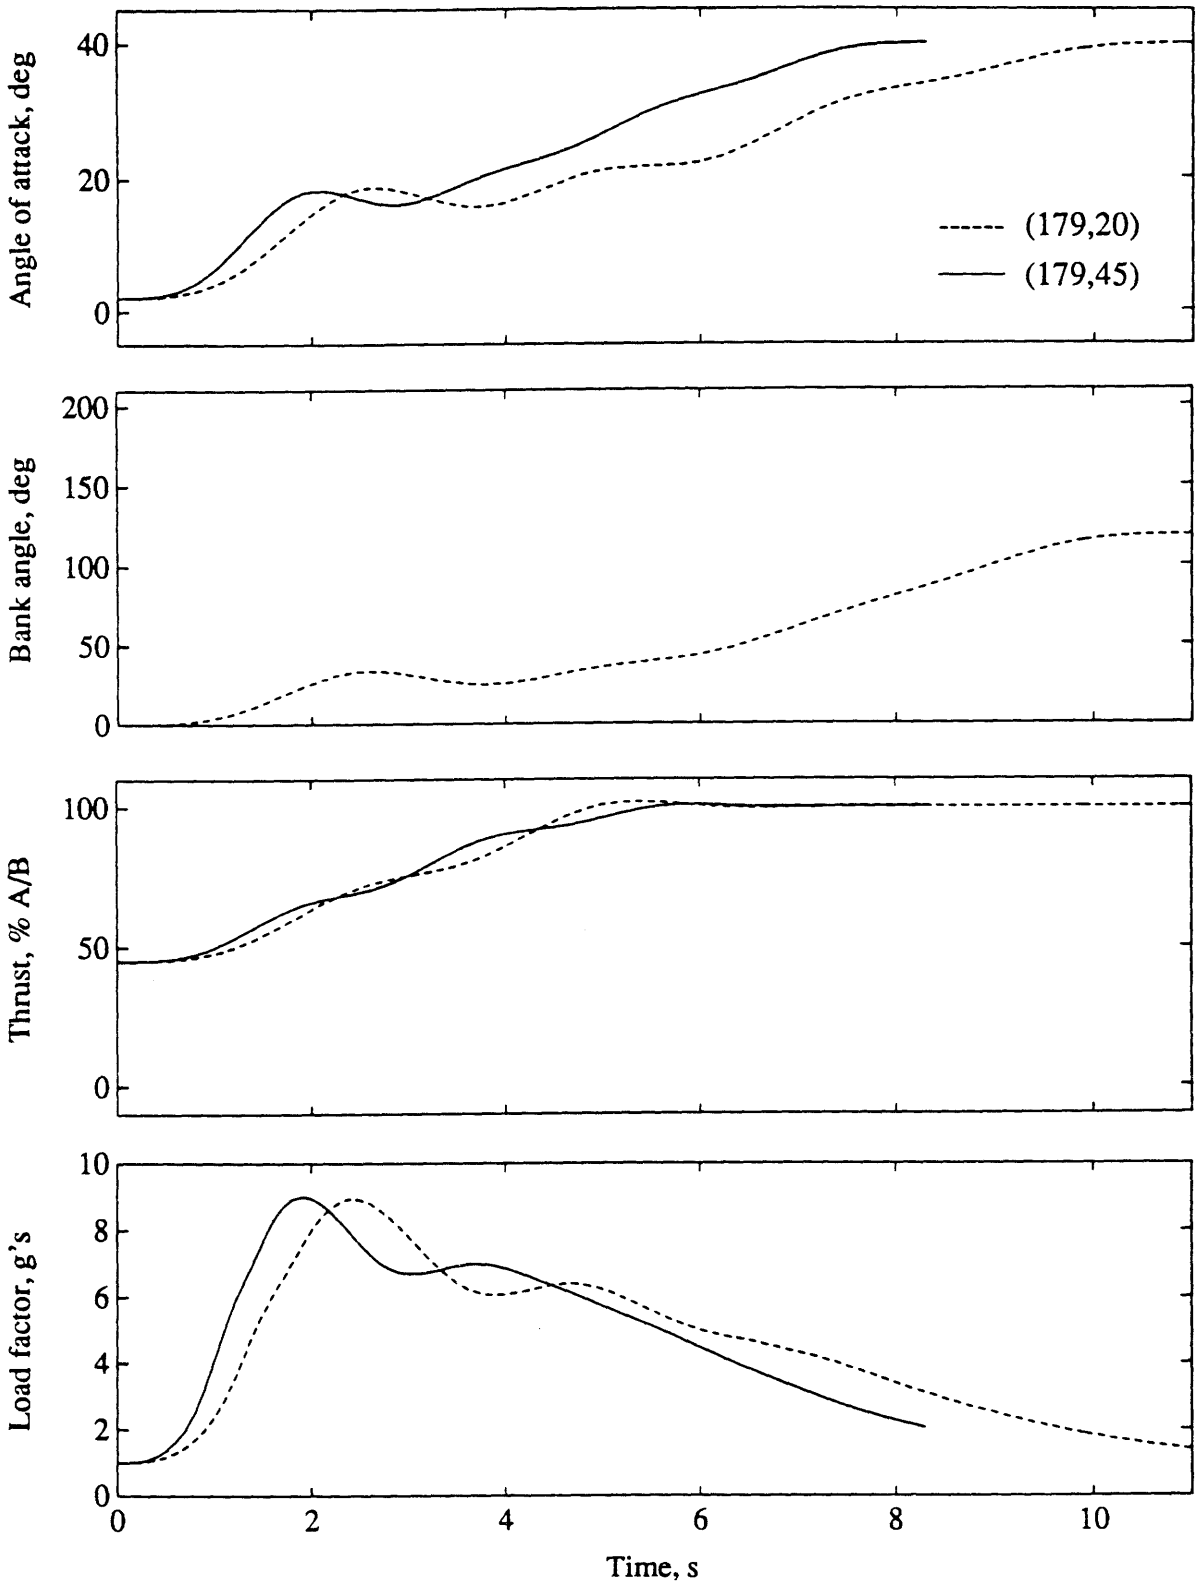

Figure 4.9. Optimal trajectories for nominal model to (179,45) and (179,20). These trajectories illustrate the switch to a split-S from a bank-and-pull. Notice the bank angle difference.

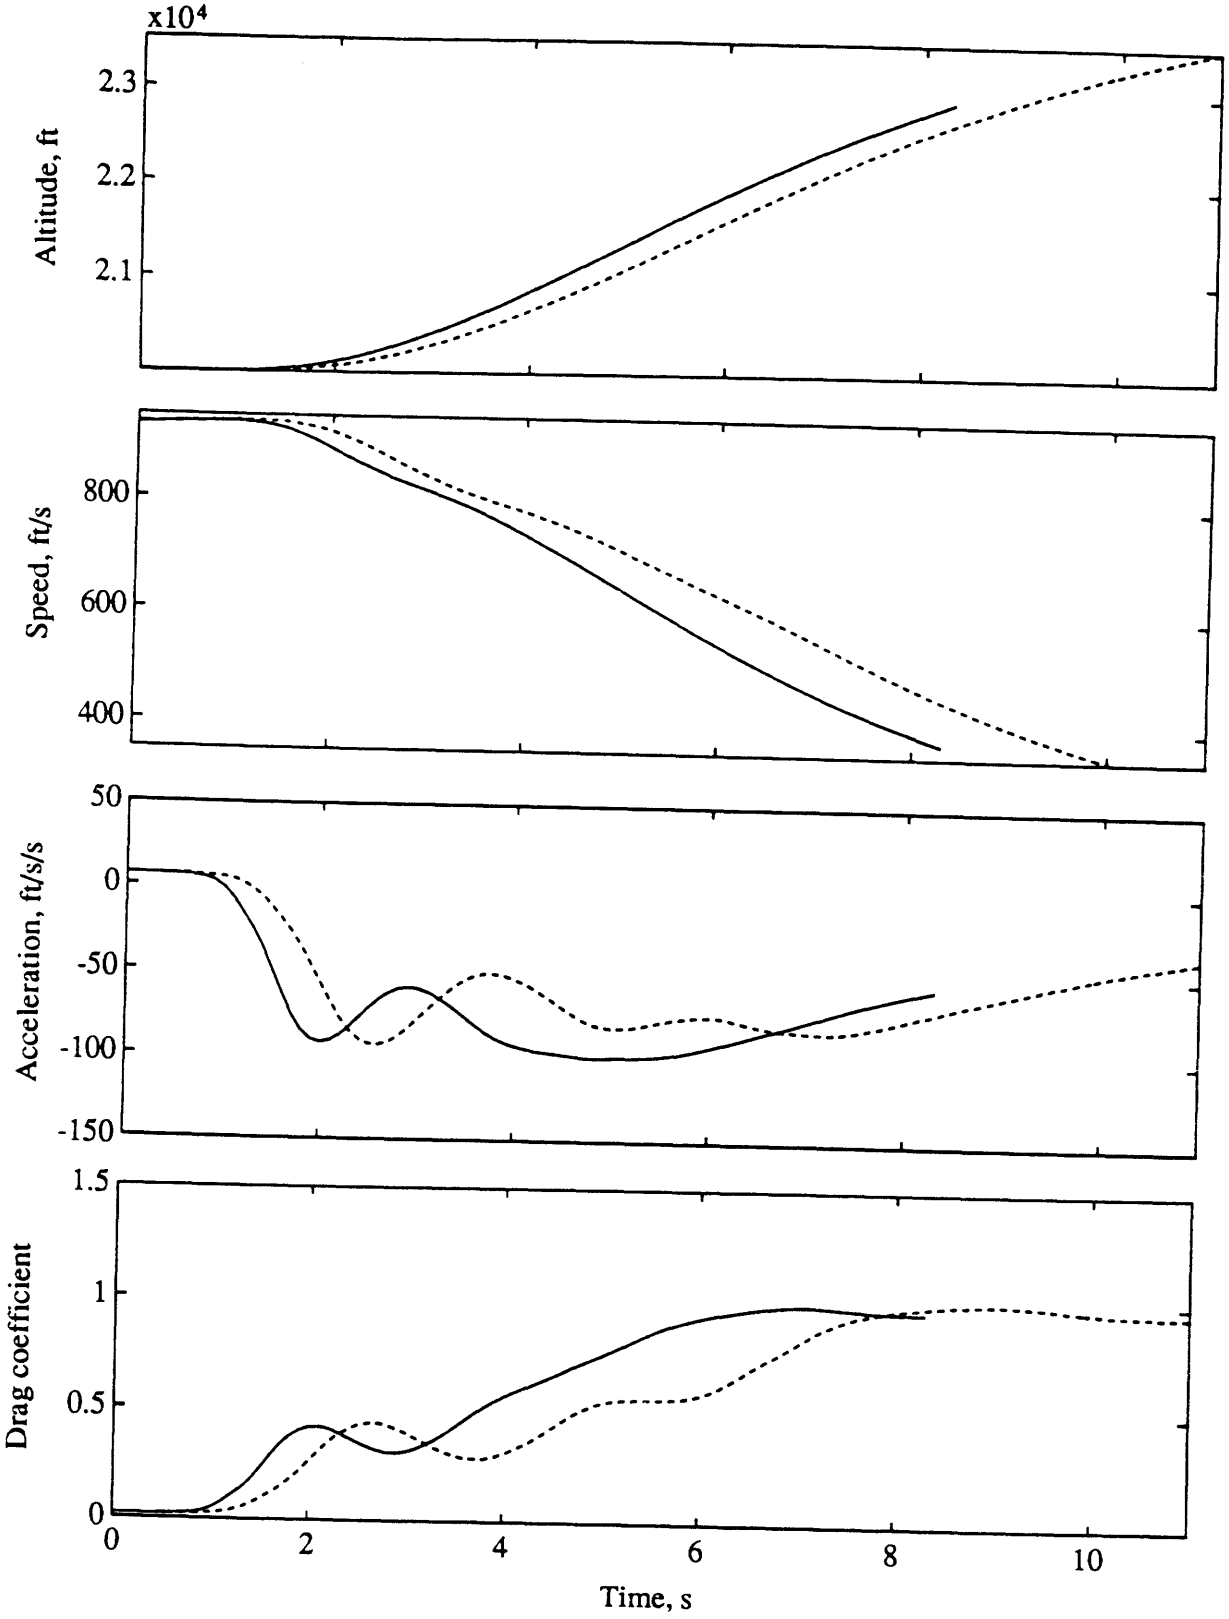

Figure 4.9 (cont'd)

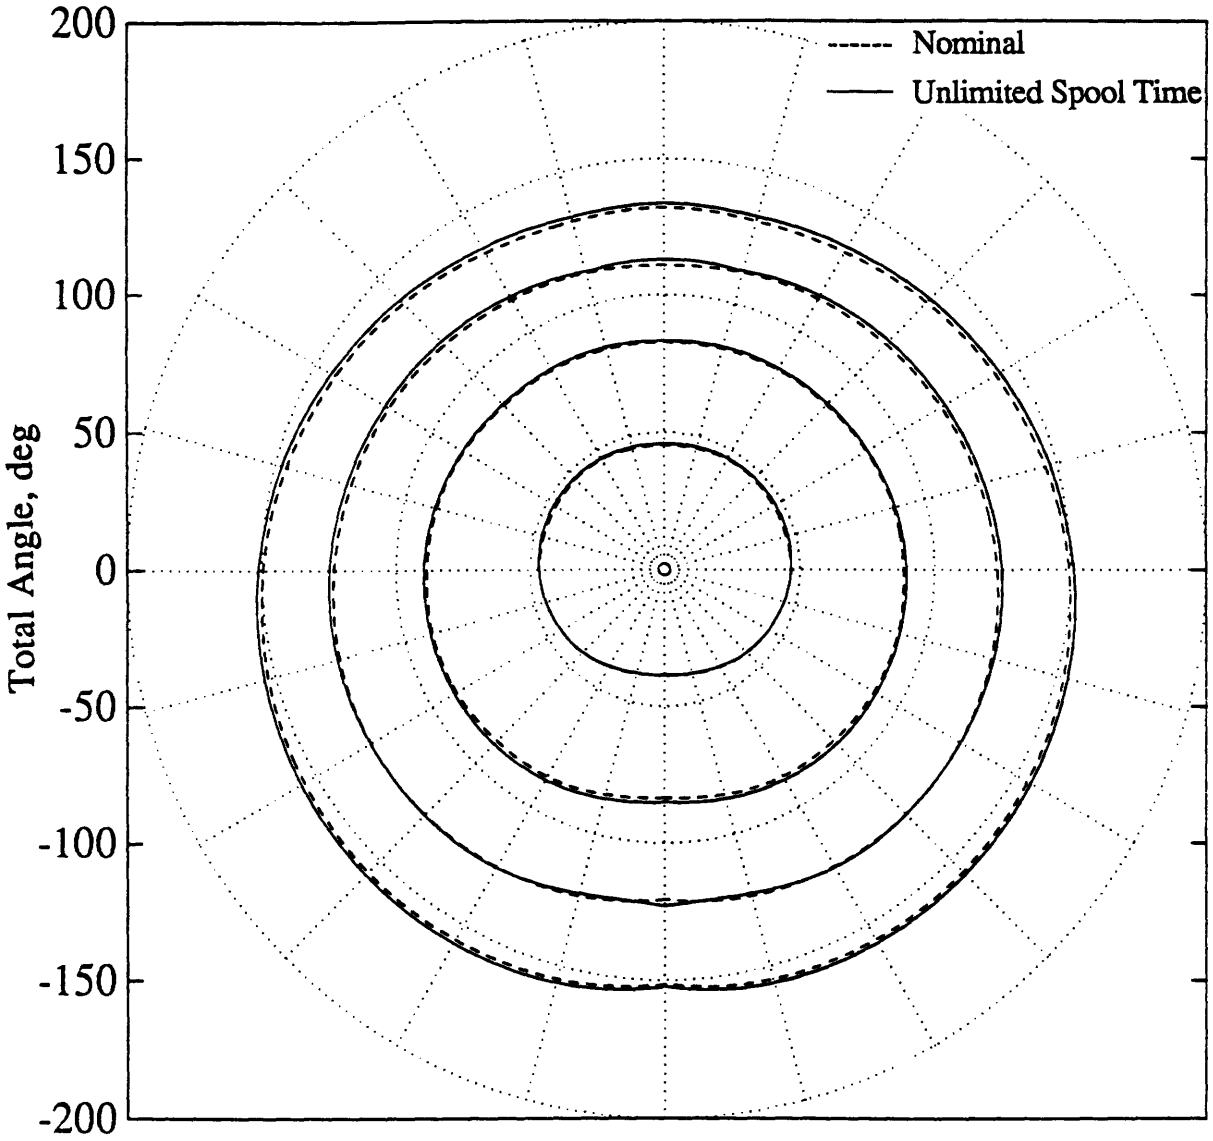

Figure 4.10. Contour plot of optimal trajectories of nominal model vs. unlimited spool time model. The contours are virtually coincident with the differences increasing as time progresses. This indicates energy management is more important for longer trajectories.

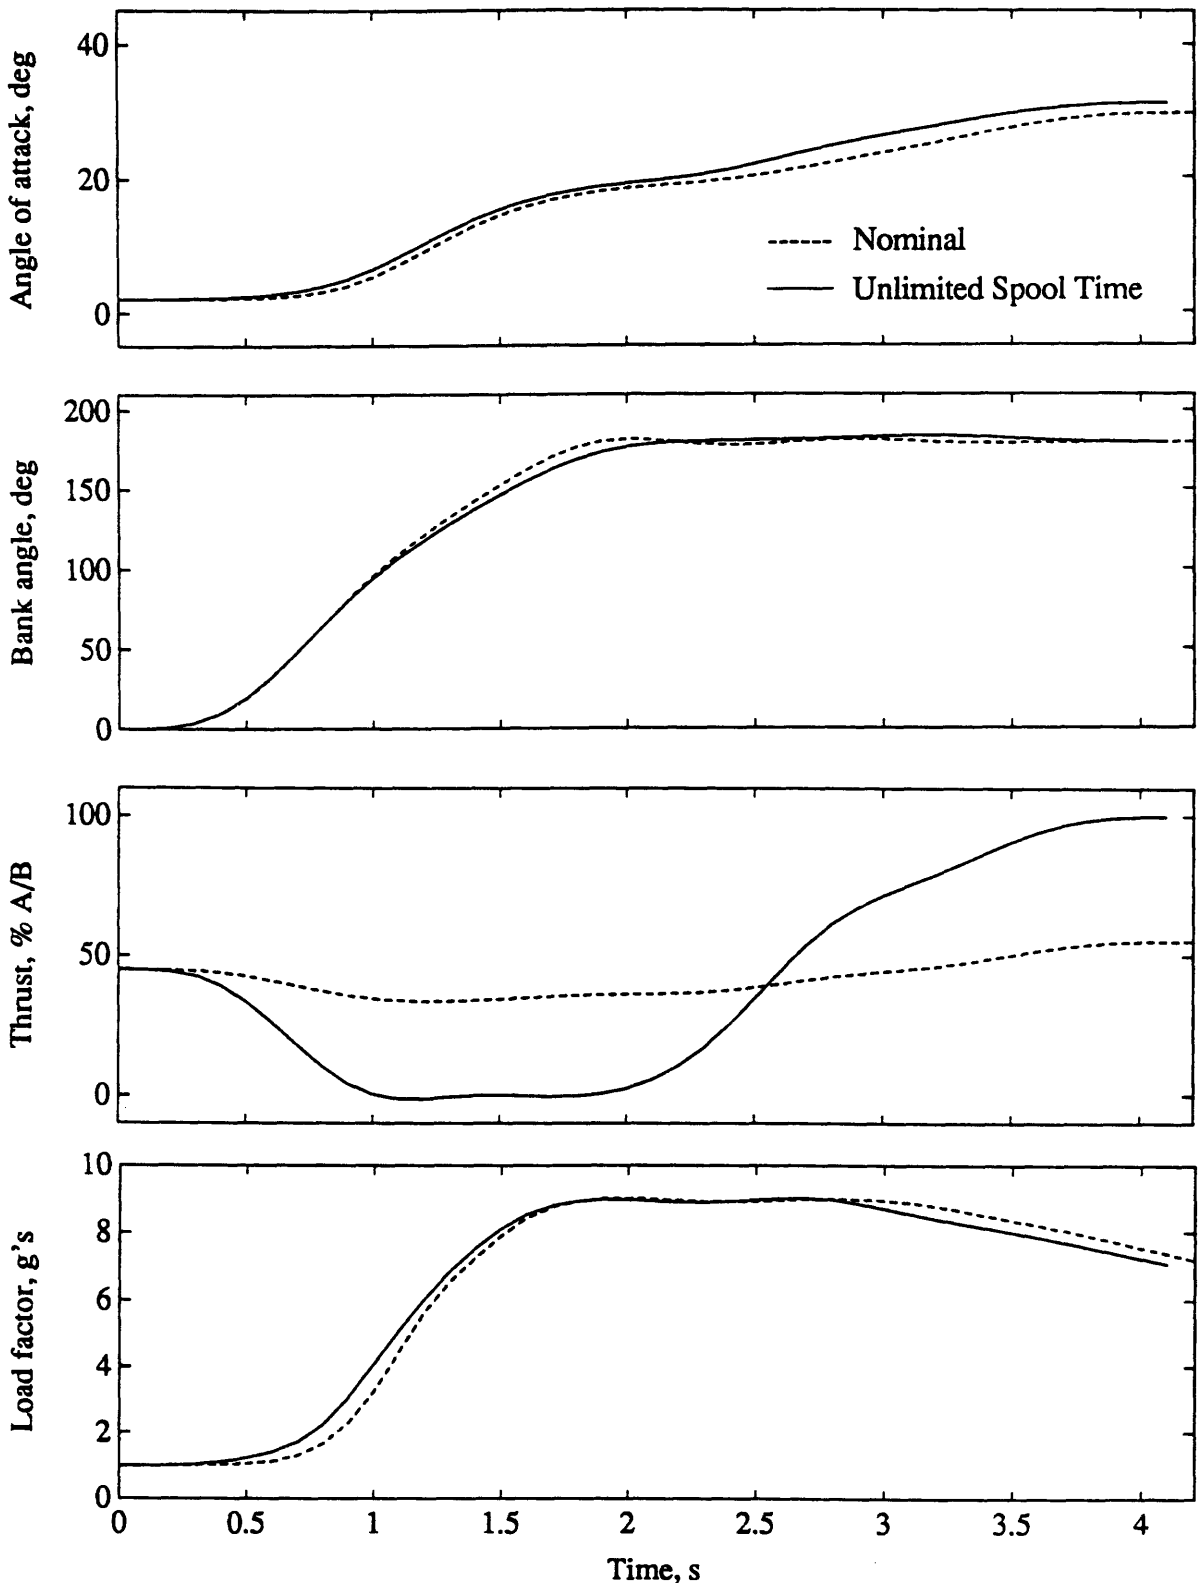

Figure 4.11. Optimal trajectories for nominal model vs. unlimited spool time model to (15, -88). Unlimited thrust history is an idle-to-A/B excursion. The spool limits of the nominal model do not allow such excursions, so the thrust stays near military power.

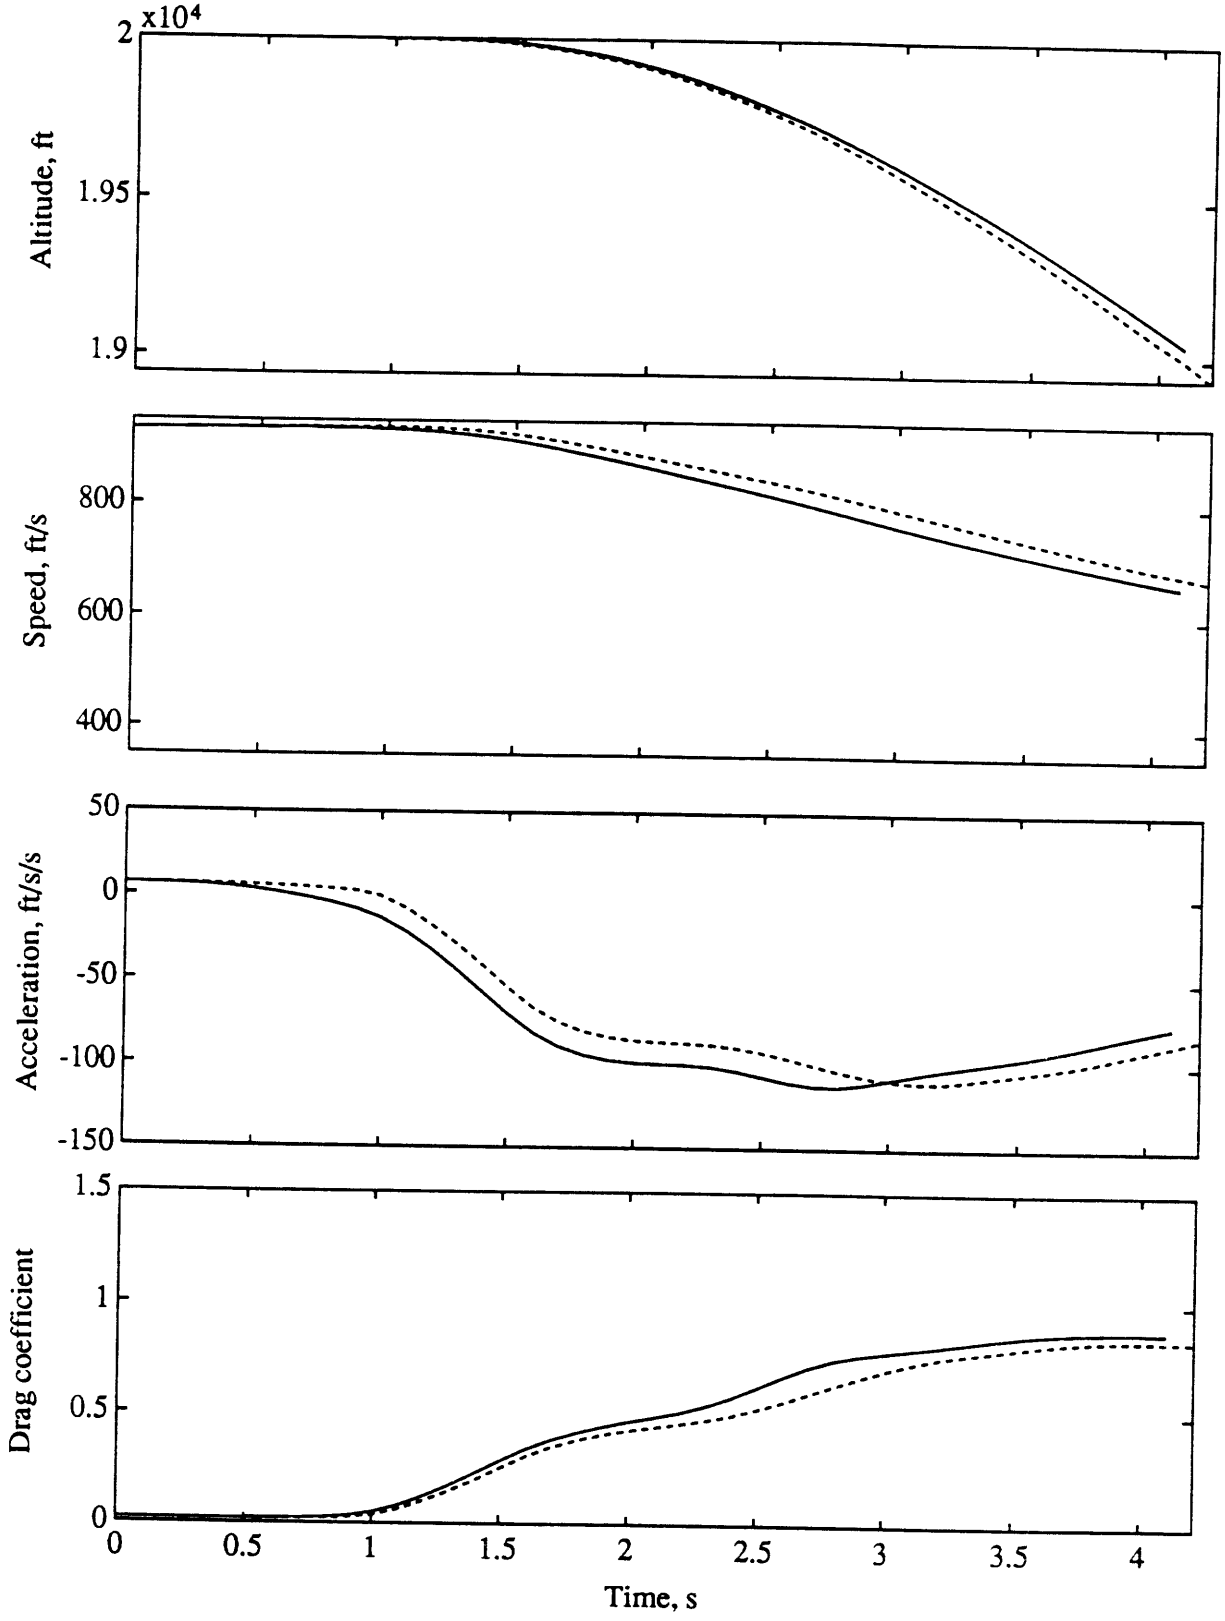

Figure 4.11 (cont'd)

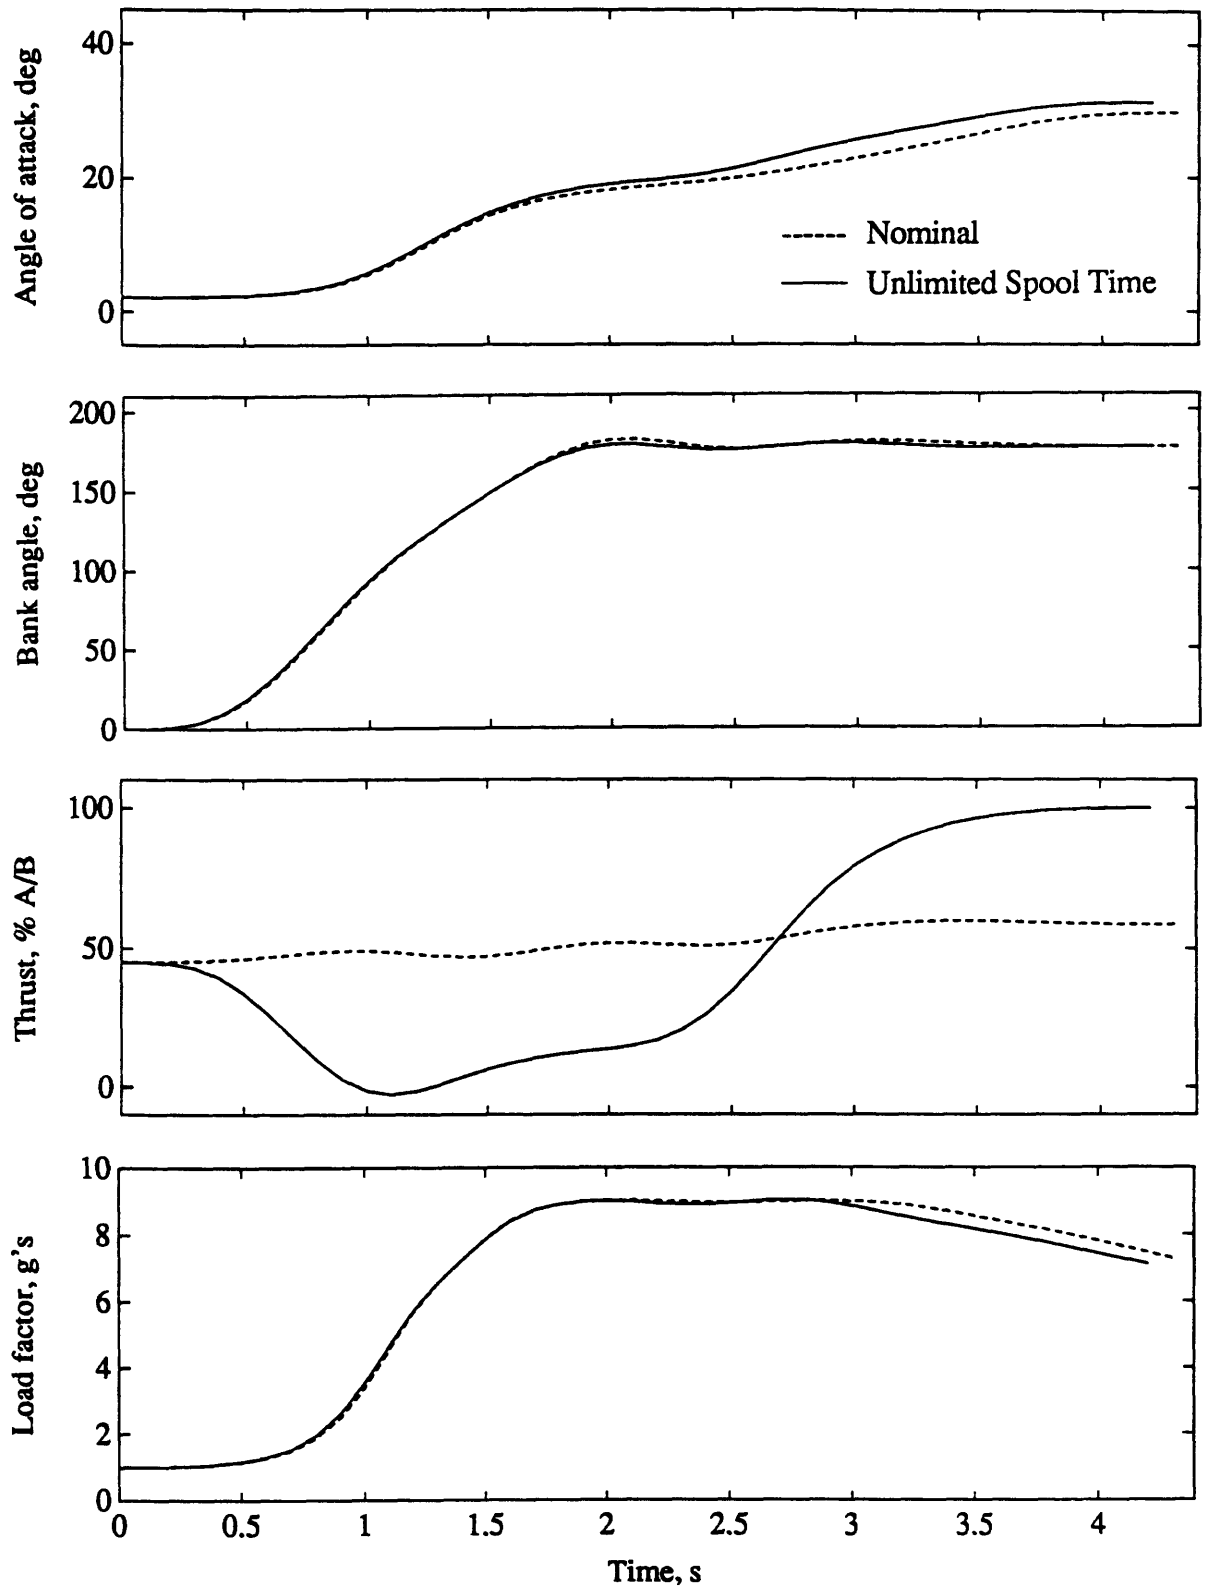

Figure 4.12. Optimal trajectories for nominal model vs. unlimited spool time model to (135, -88). The unlimited thrust strategy does not change for straight-down targets regardless of heading.

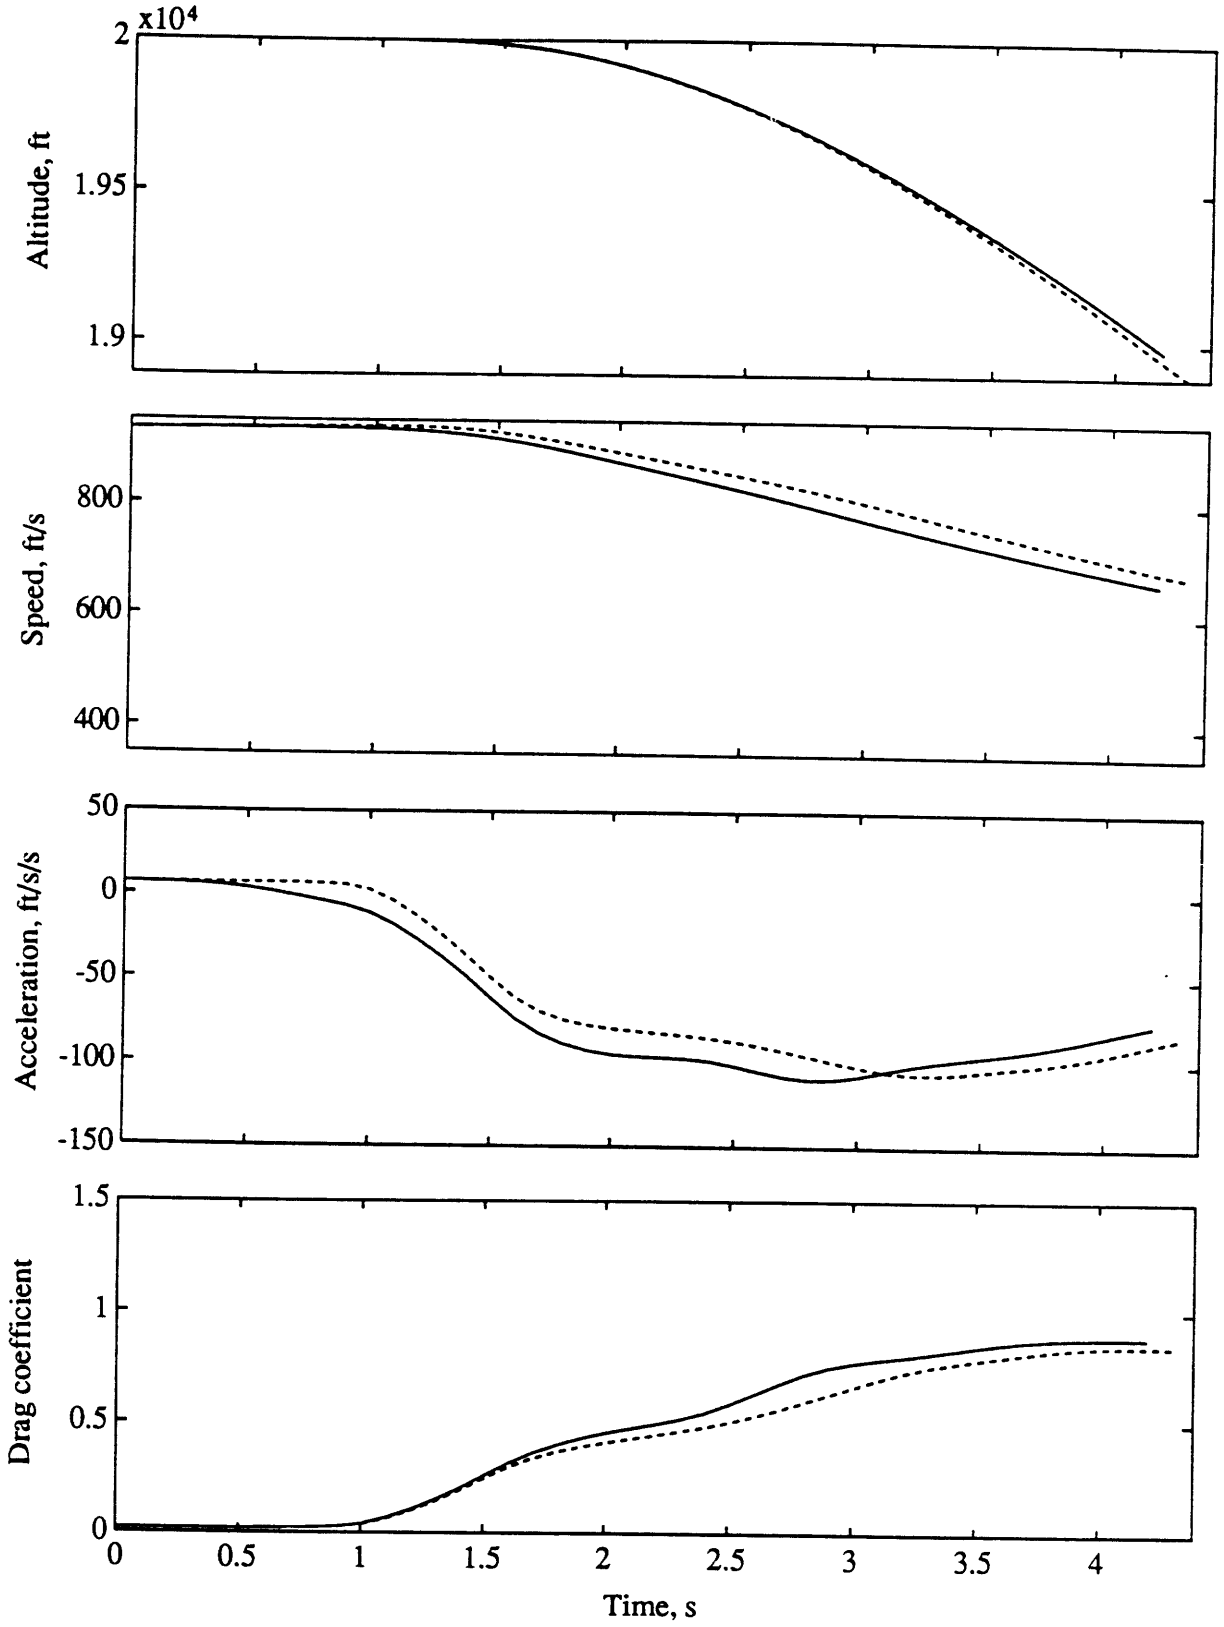

Figure 4.12 (cont'd)

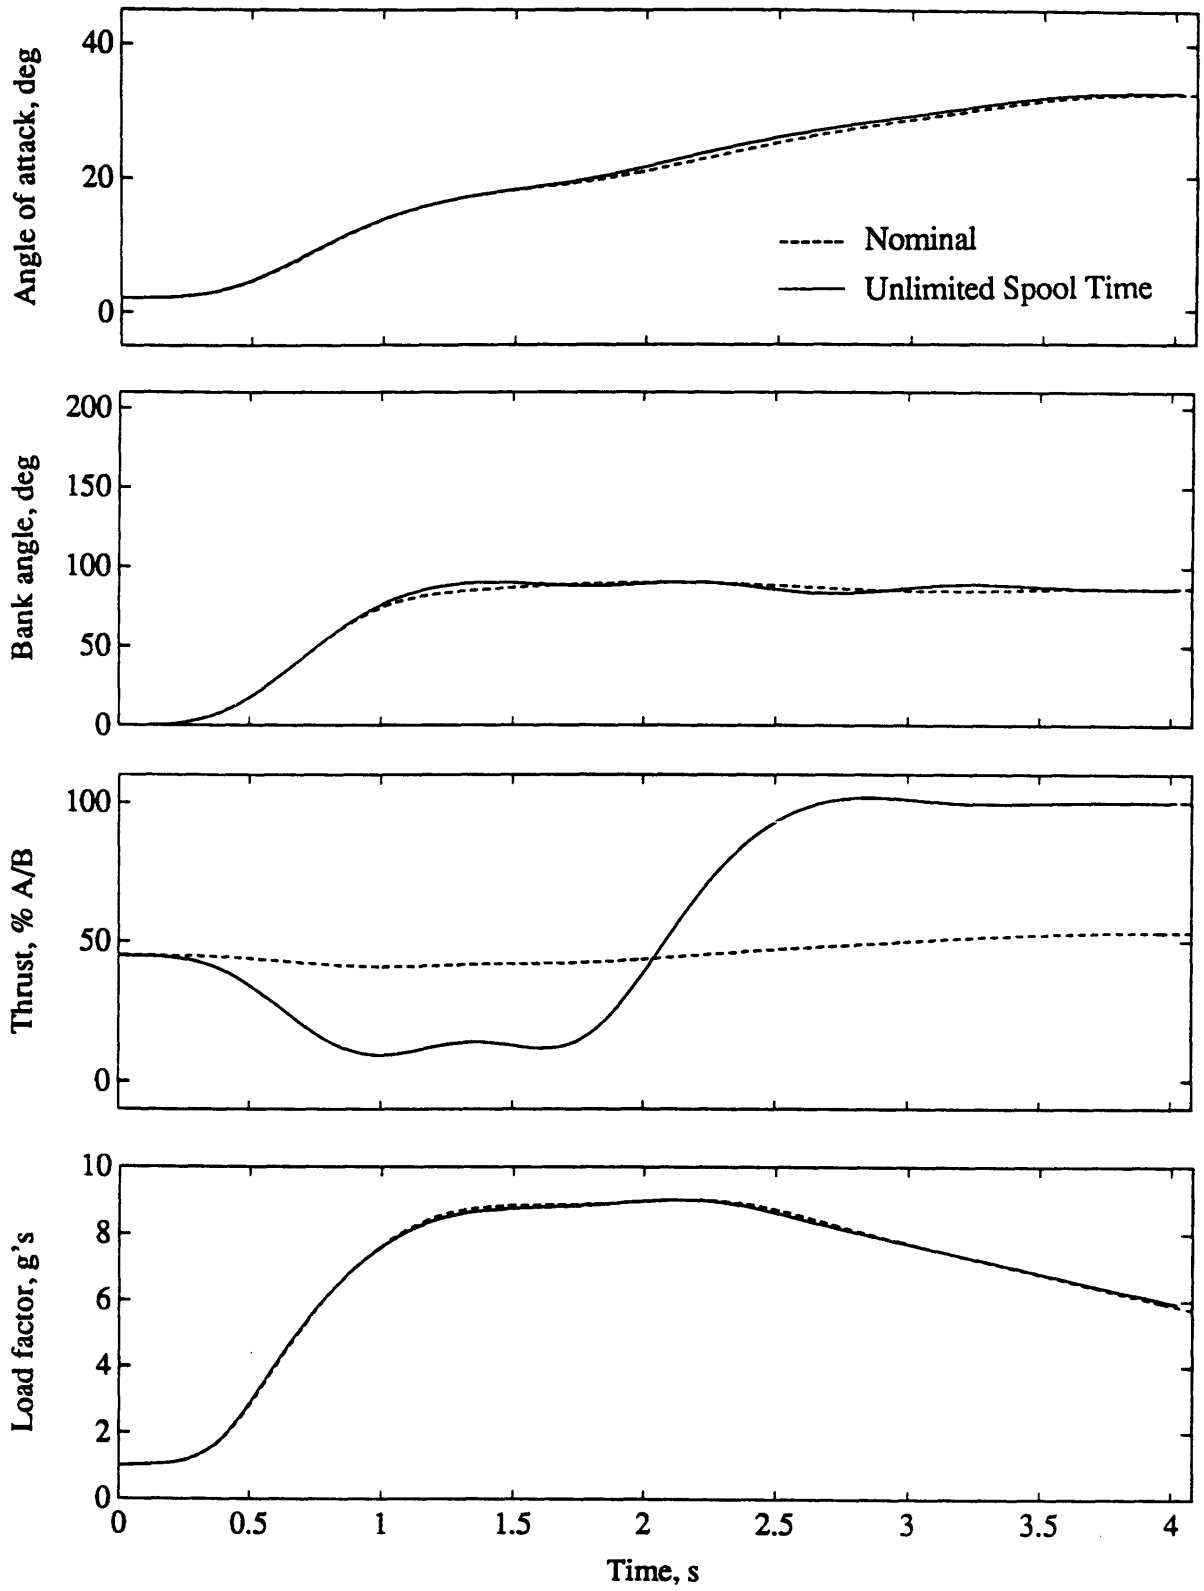

Figure 4.13. Optimal trajectories for nominal model vs. unlimited spool time model to (90,0). The idle-to-A/B excursion is also present in trajectories of moderate duration.

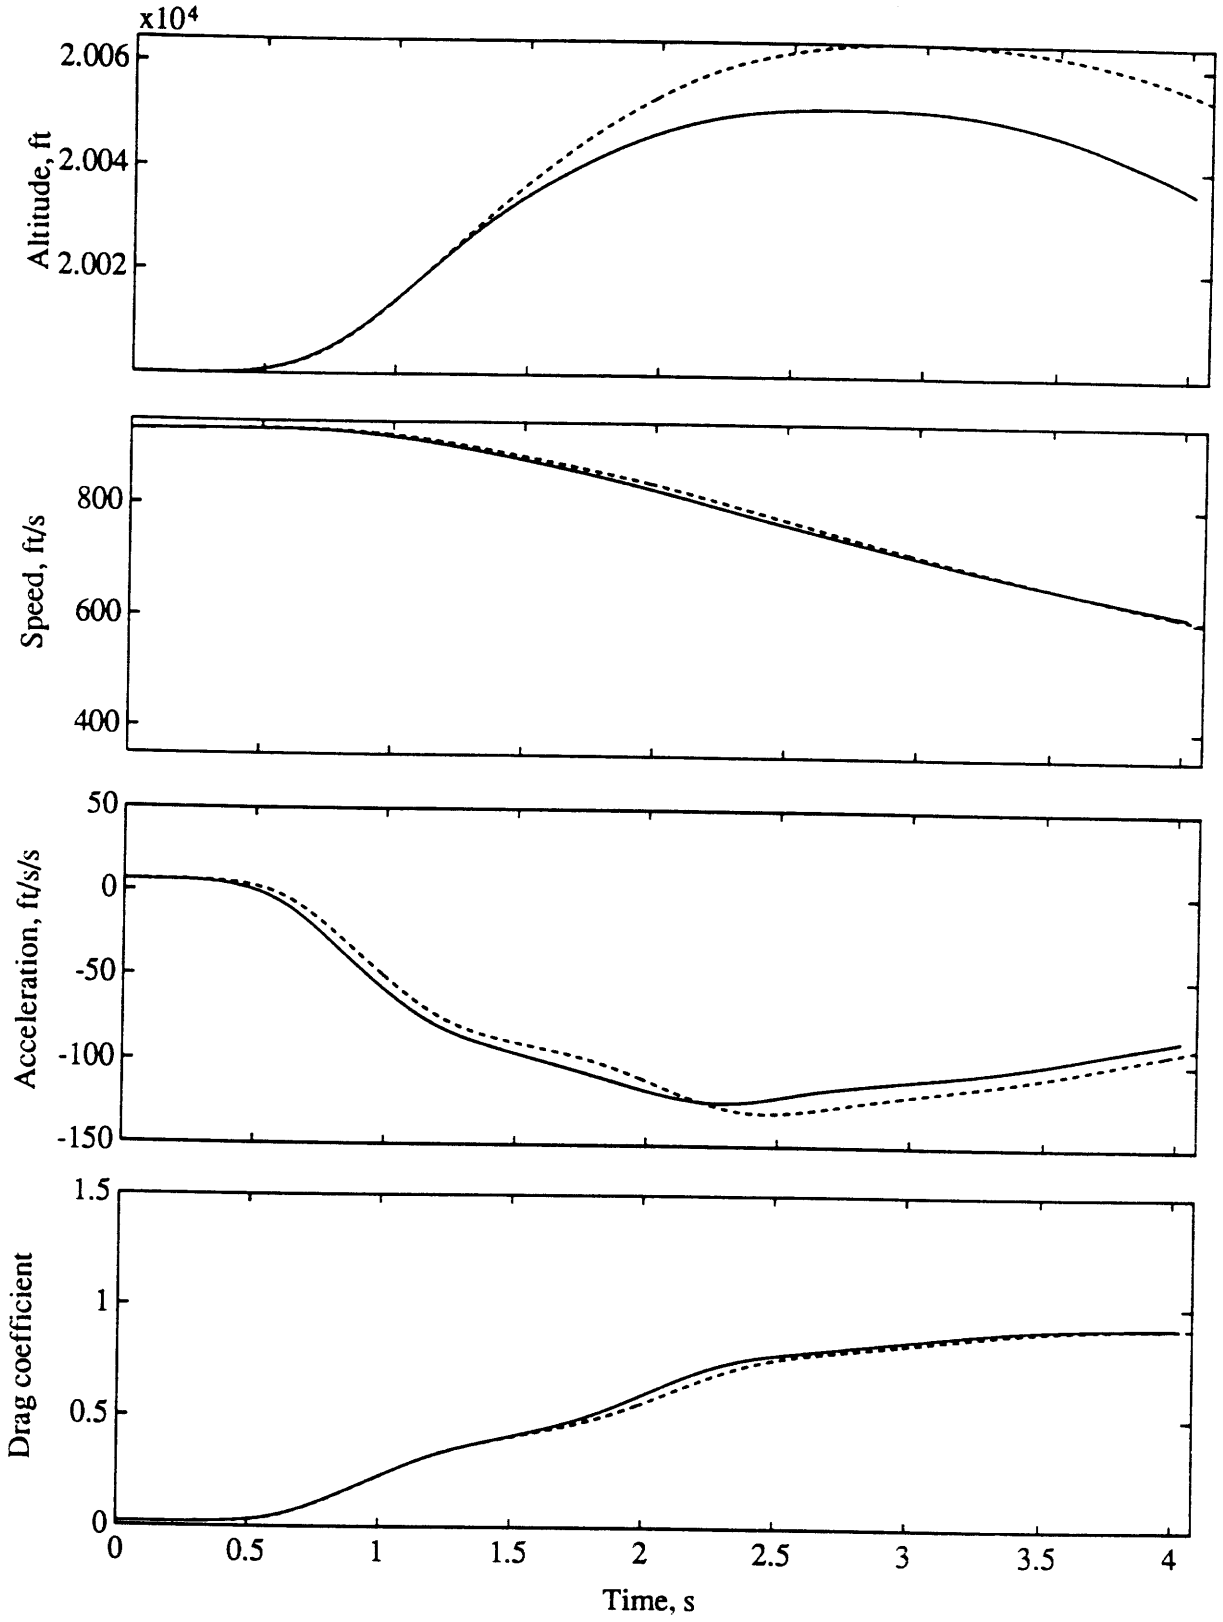

Figure 4.13 (cont'd)

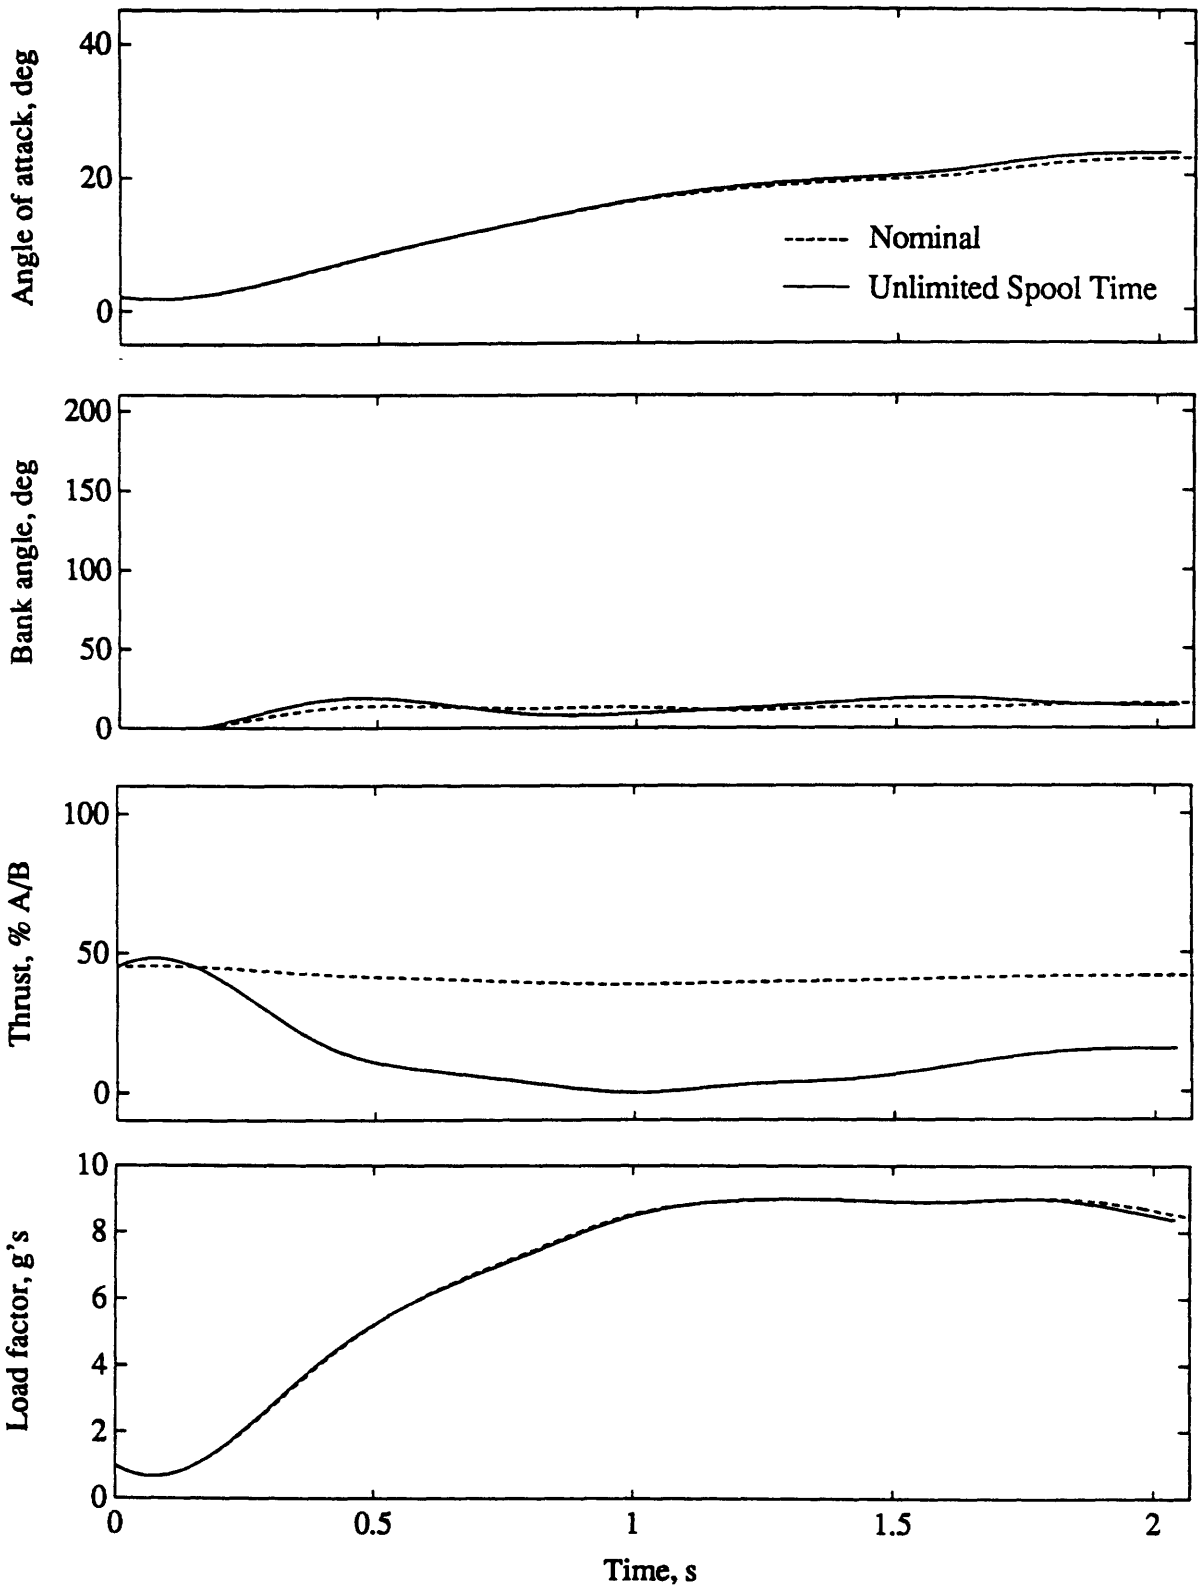

Figure 4.14. Optimal trajectories for nominal model vs. unlimited spool time model to (15, 45). Nominal thrust is near idle for most of the trajectory because the bleed rate is less and g_{max} can be maintained.

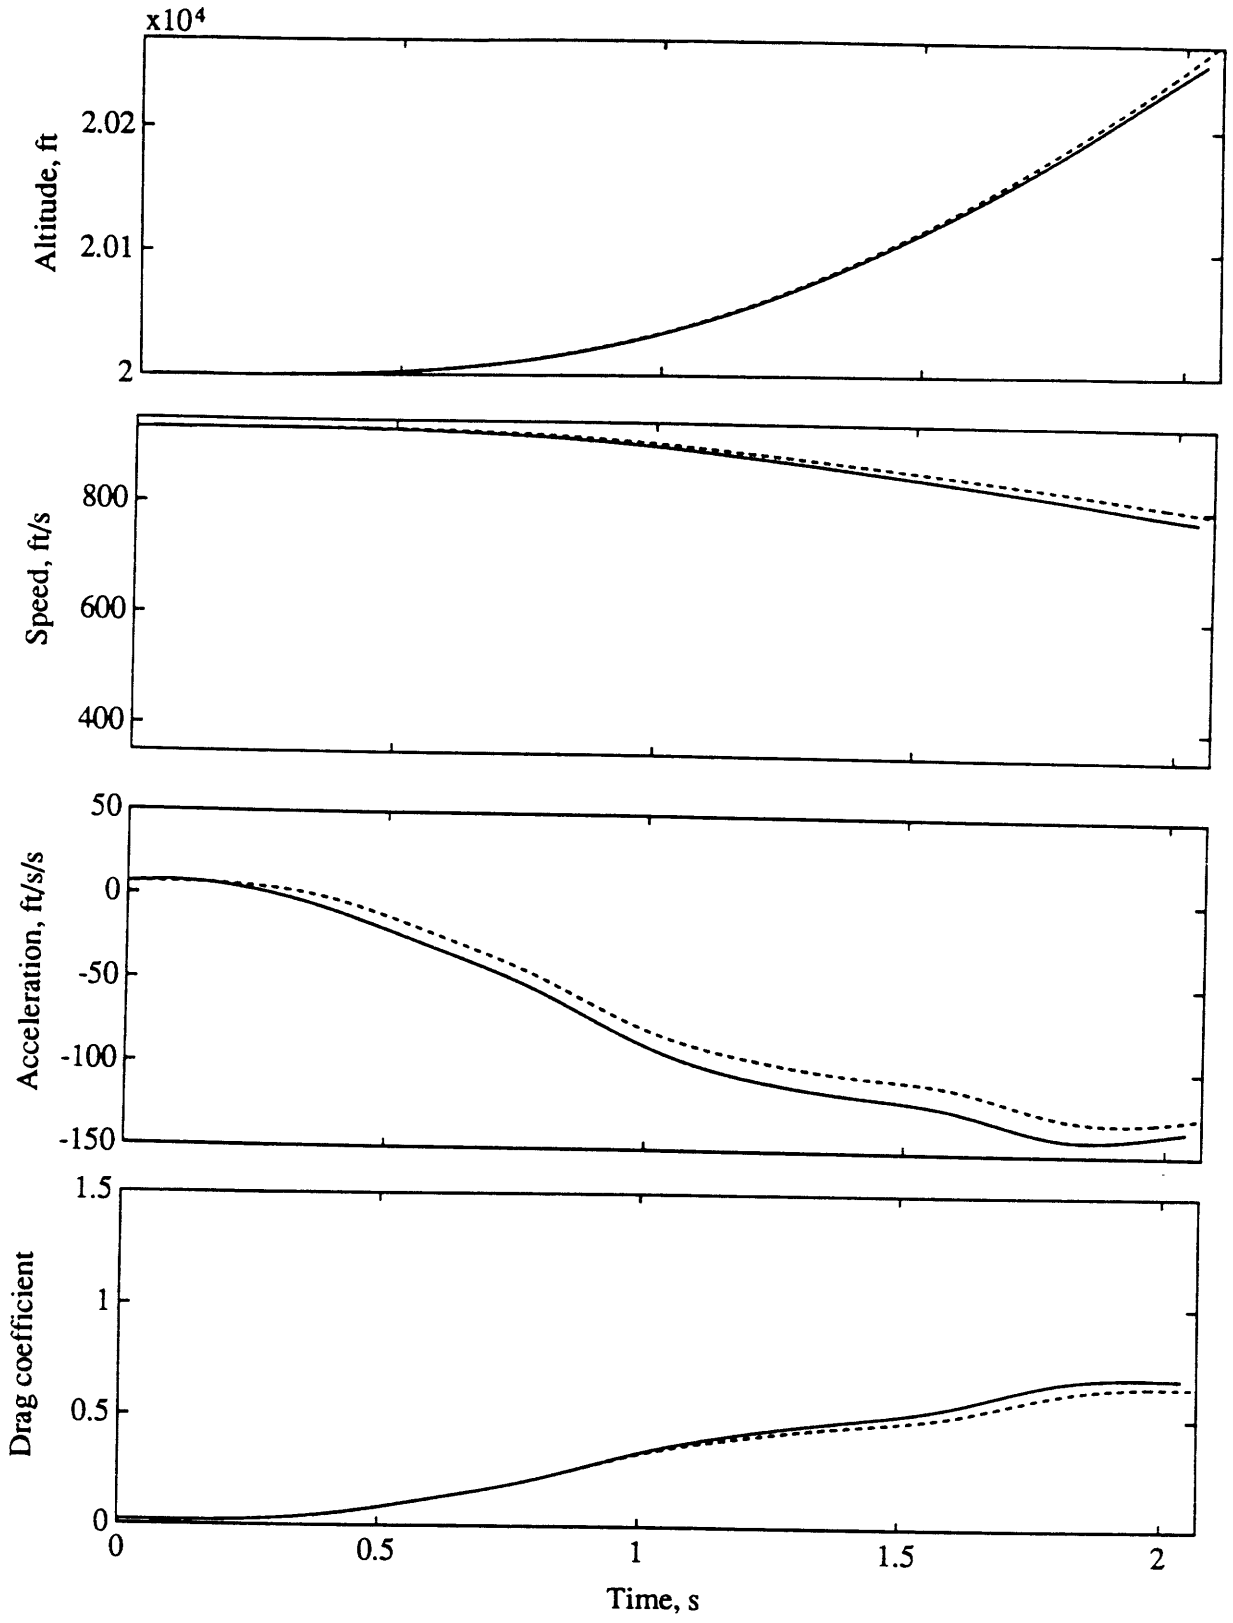

Figure 4.14 (cont'd)

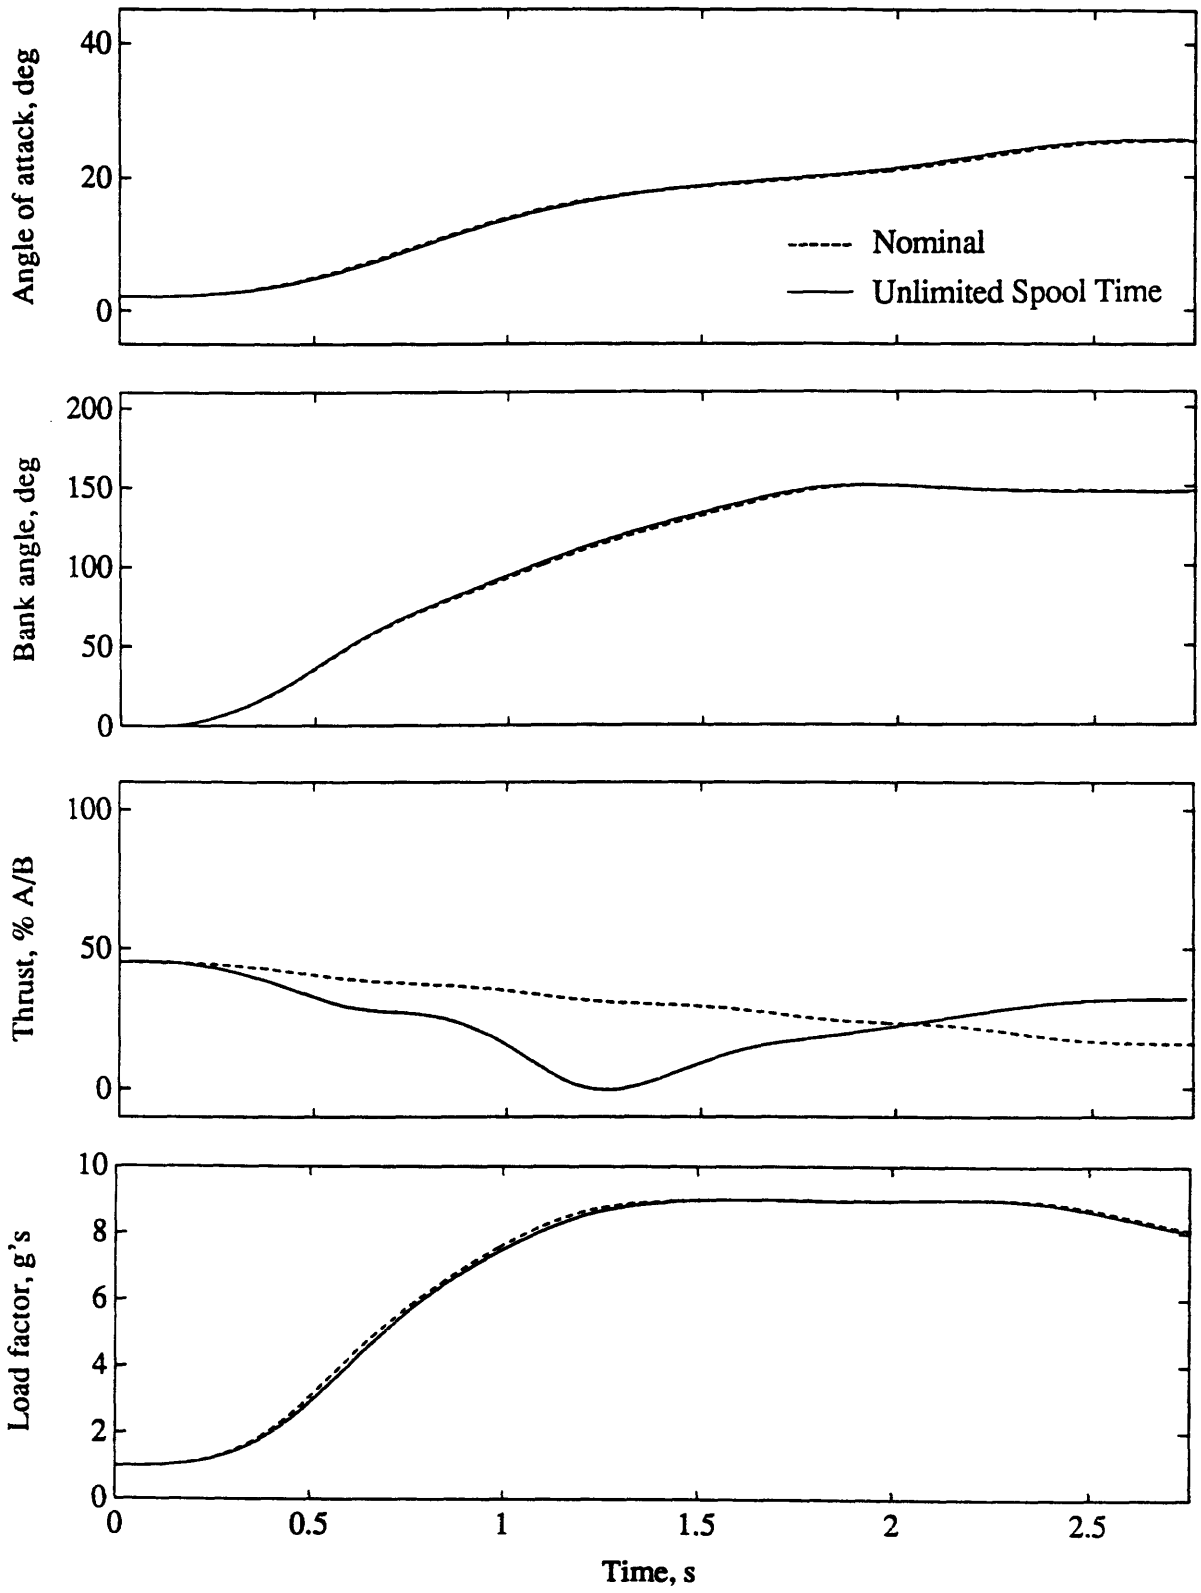

Figure 4.15. Optimal trajectories for nominal model vs. unlimited spool time model to (45,-45). Bleed rate is higher than for (15,45), so the nominal thrust starts back up to military power.

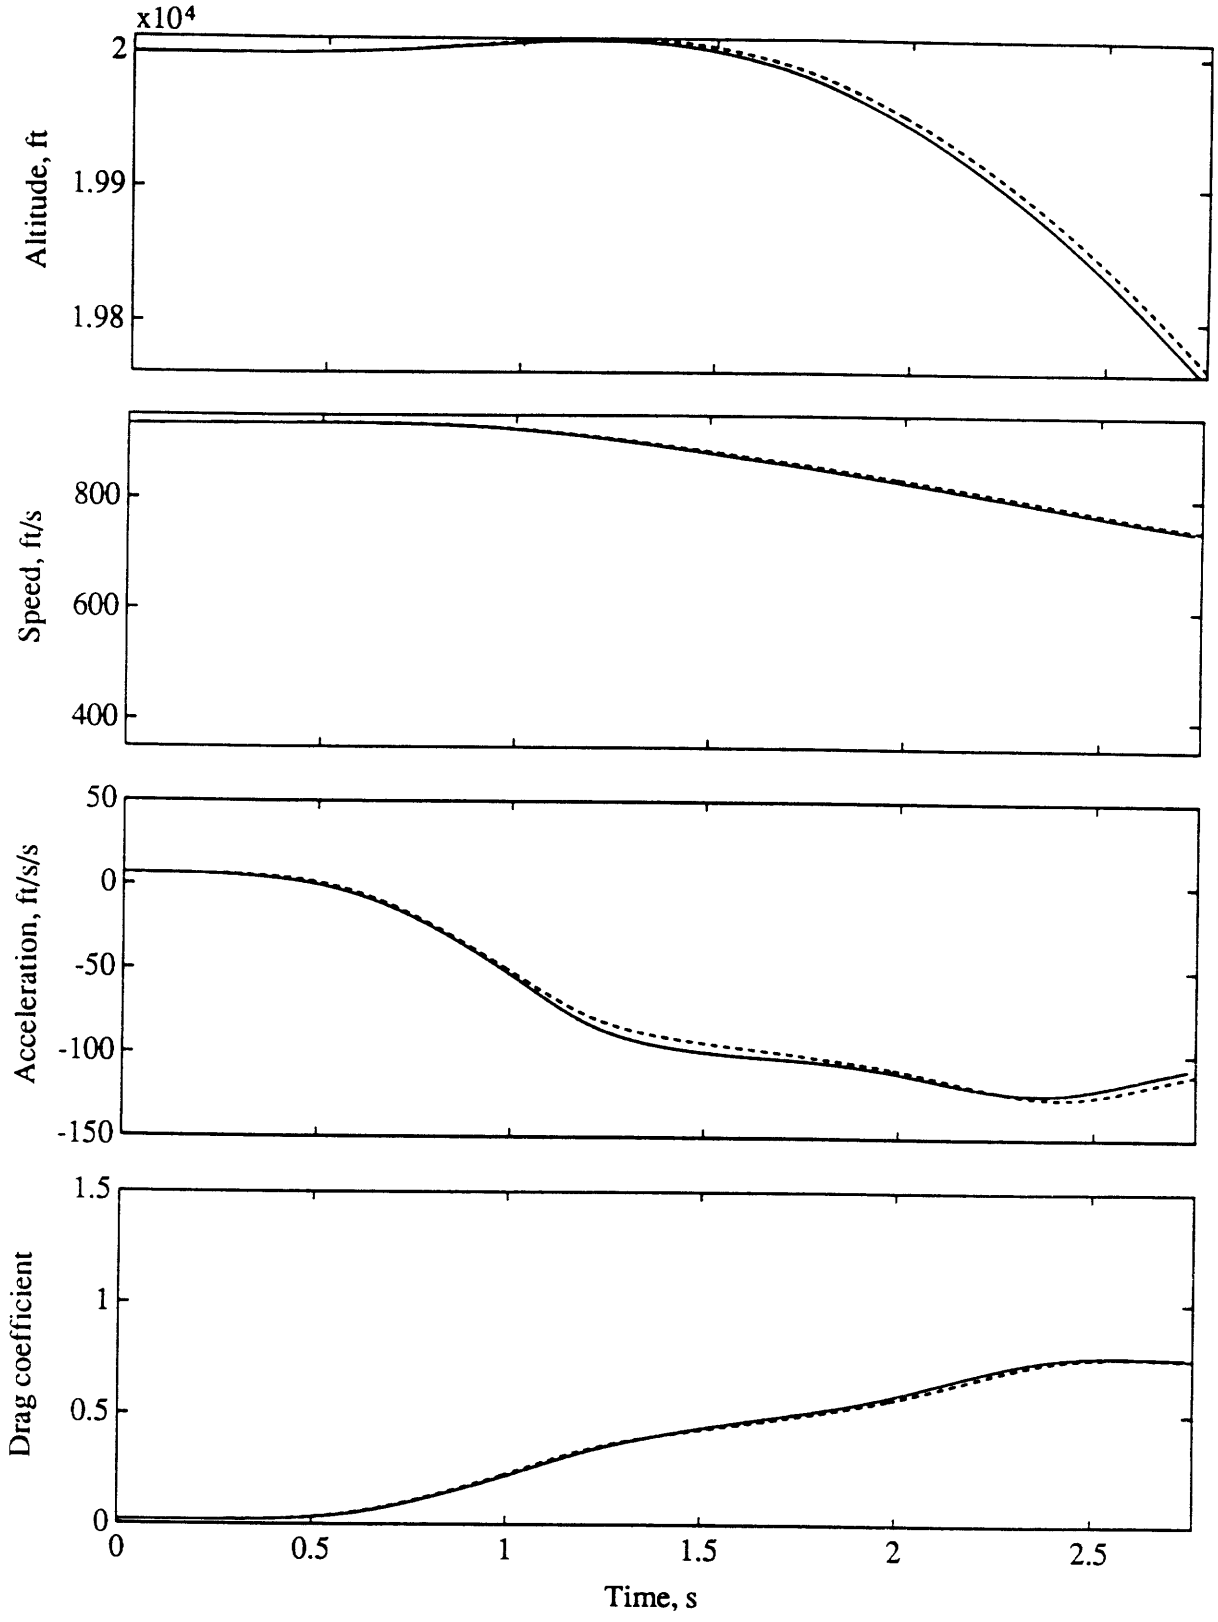

Figure 4.15 (cont'd)

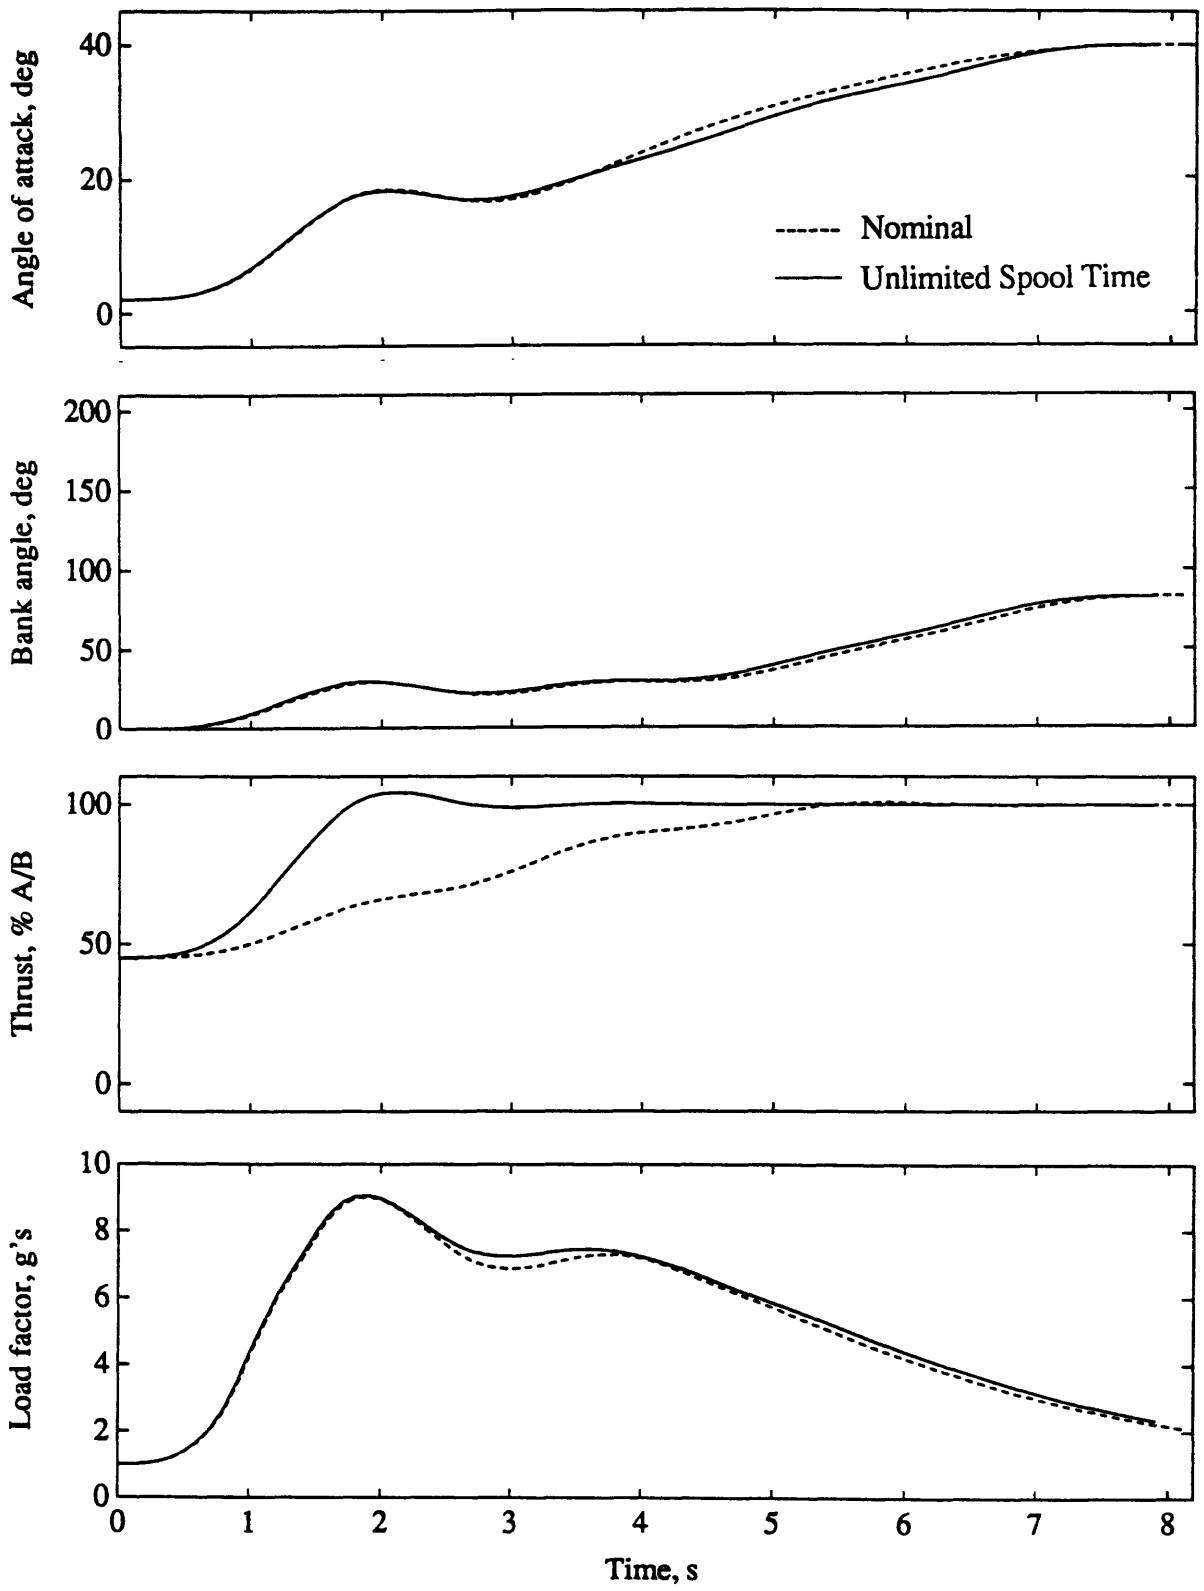

Figure 4.16. Optimal trajectories for nominal model vs. unlimited spool time model to (165,45). For this long trajectory, angle of attack reaches maximum and drag is high. The speed loss is correspondingly high, so the thrust must immediately go to A/B.

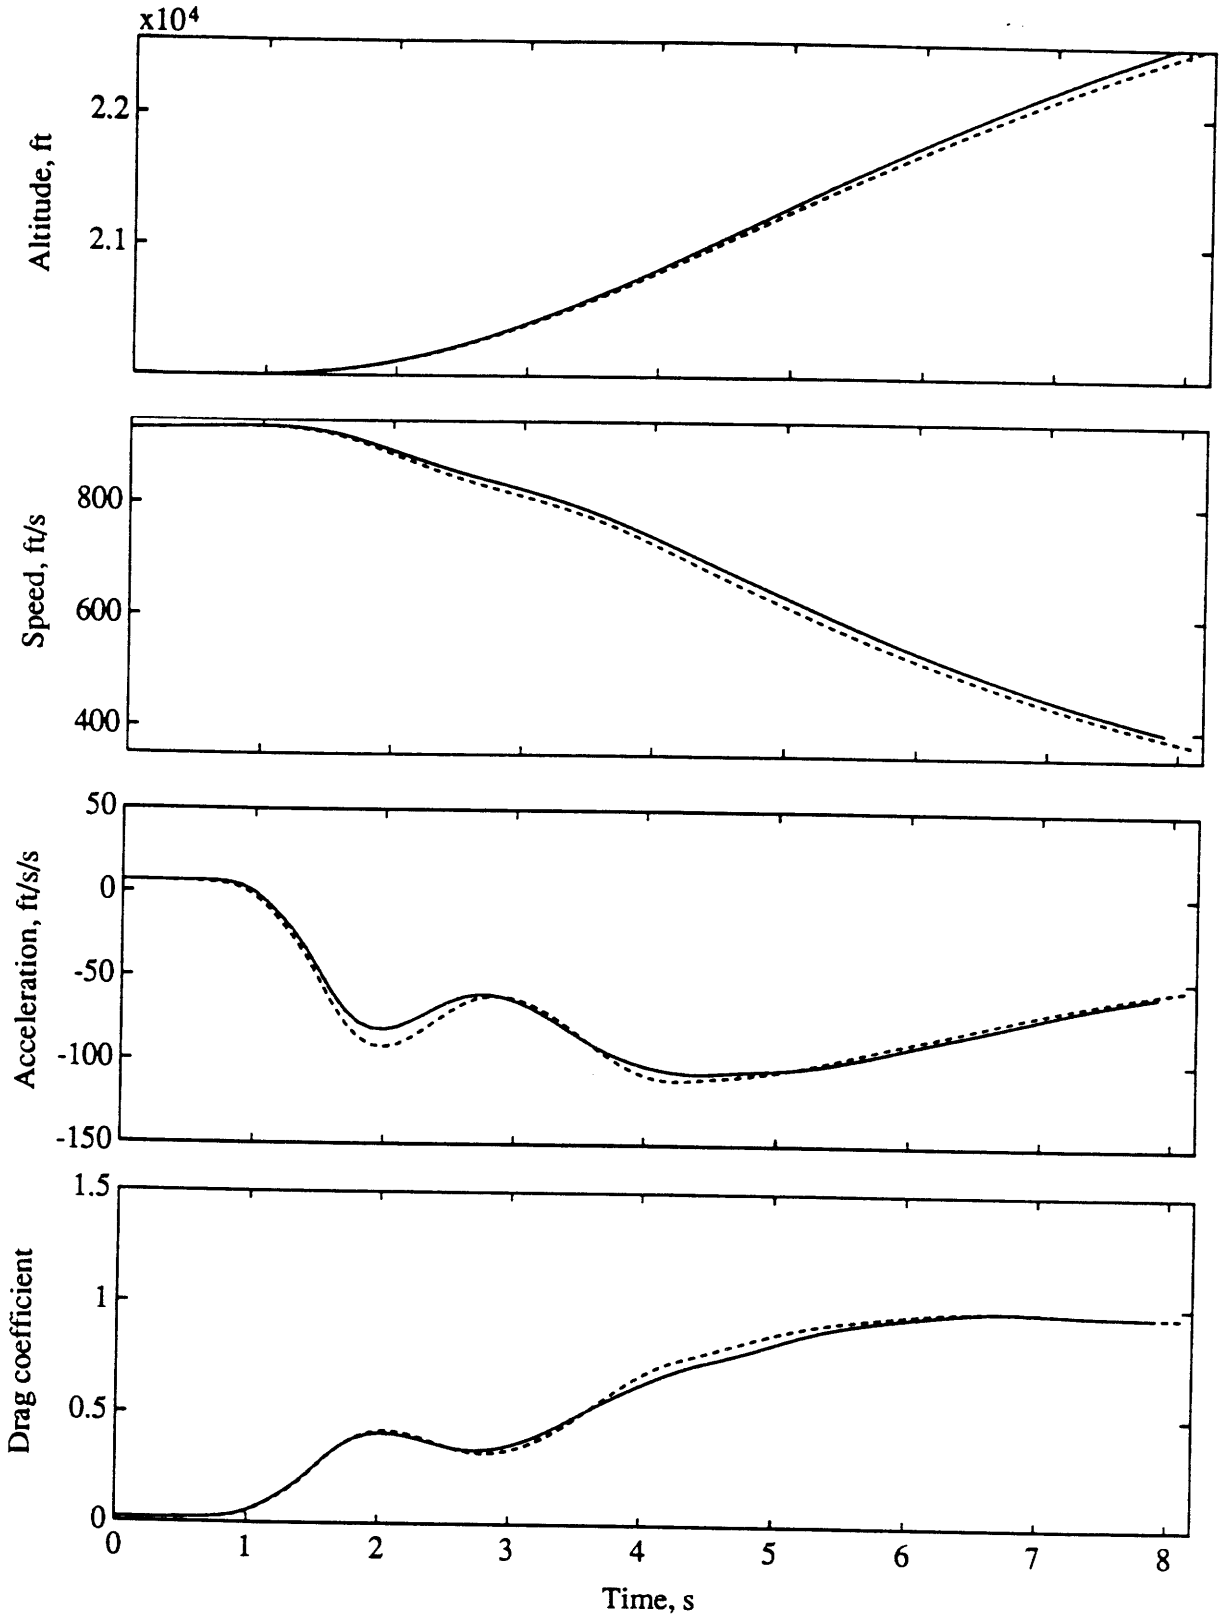

Figure 4.16 (cont'd)

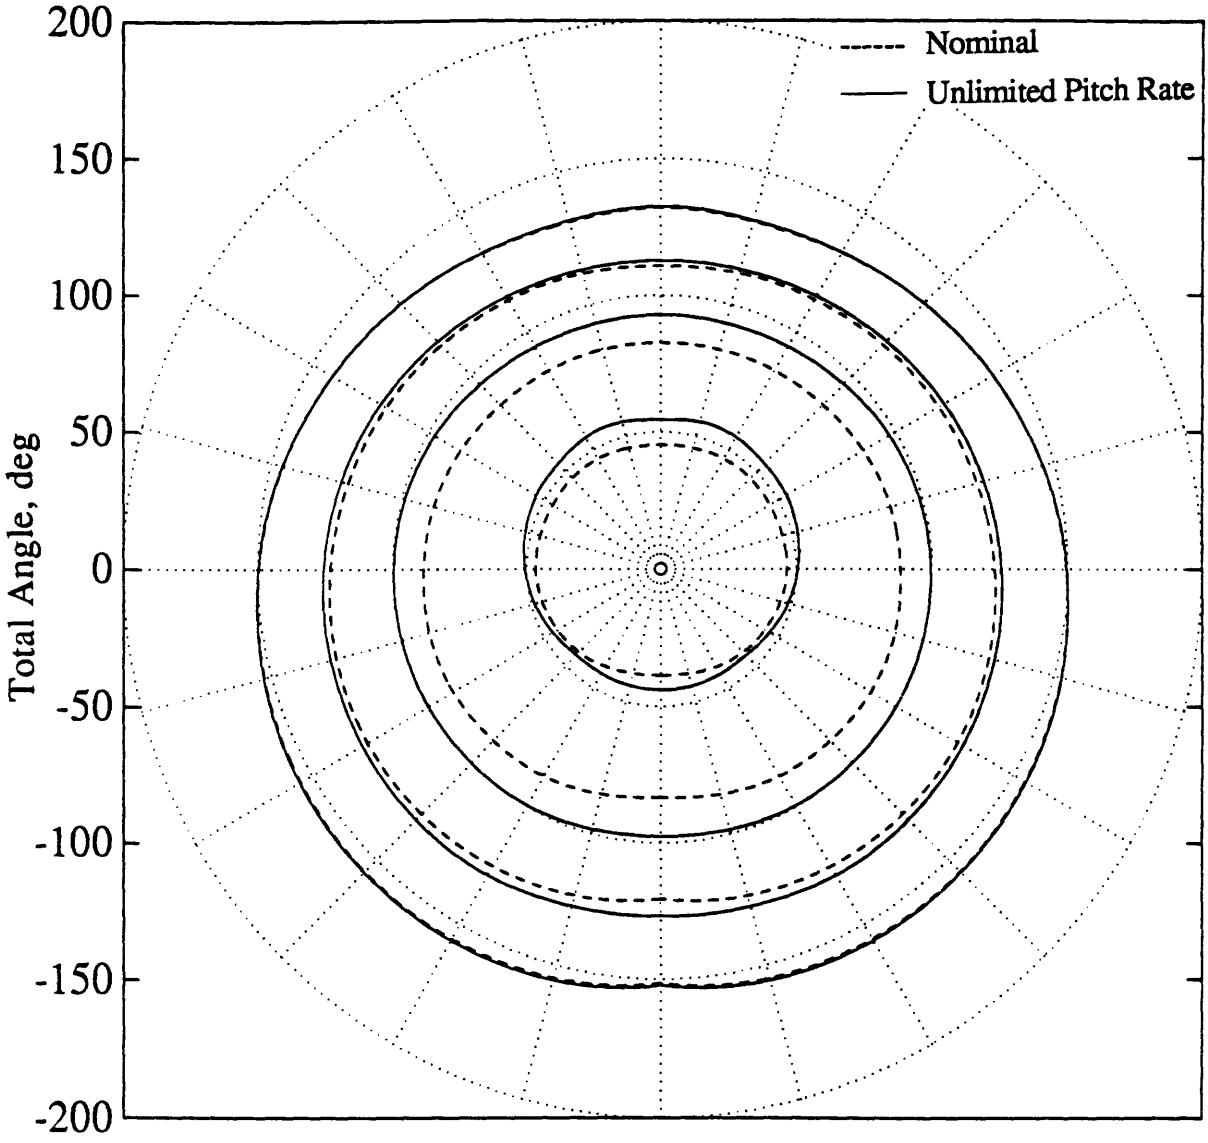

Figure 4.17. Contour plot of optimal trajectories of nominal model vs. unlimited pitch rate model. The differences are strong in the 2 and 4 s contours because high angles of attack do not have much time to bleed the airspeed significantly. By 8 s, energy management becomes more important; and the unlimited model cannot jump immediately to high α because of the associated high drag.

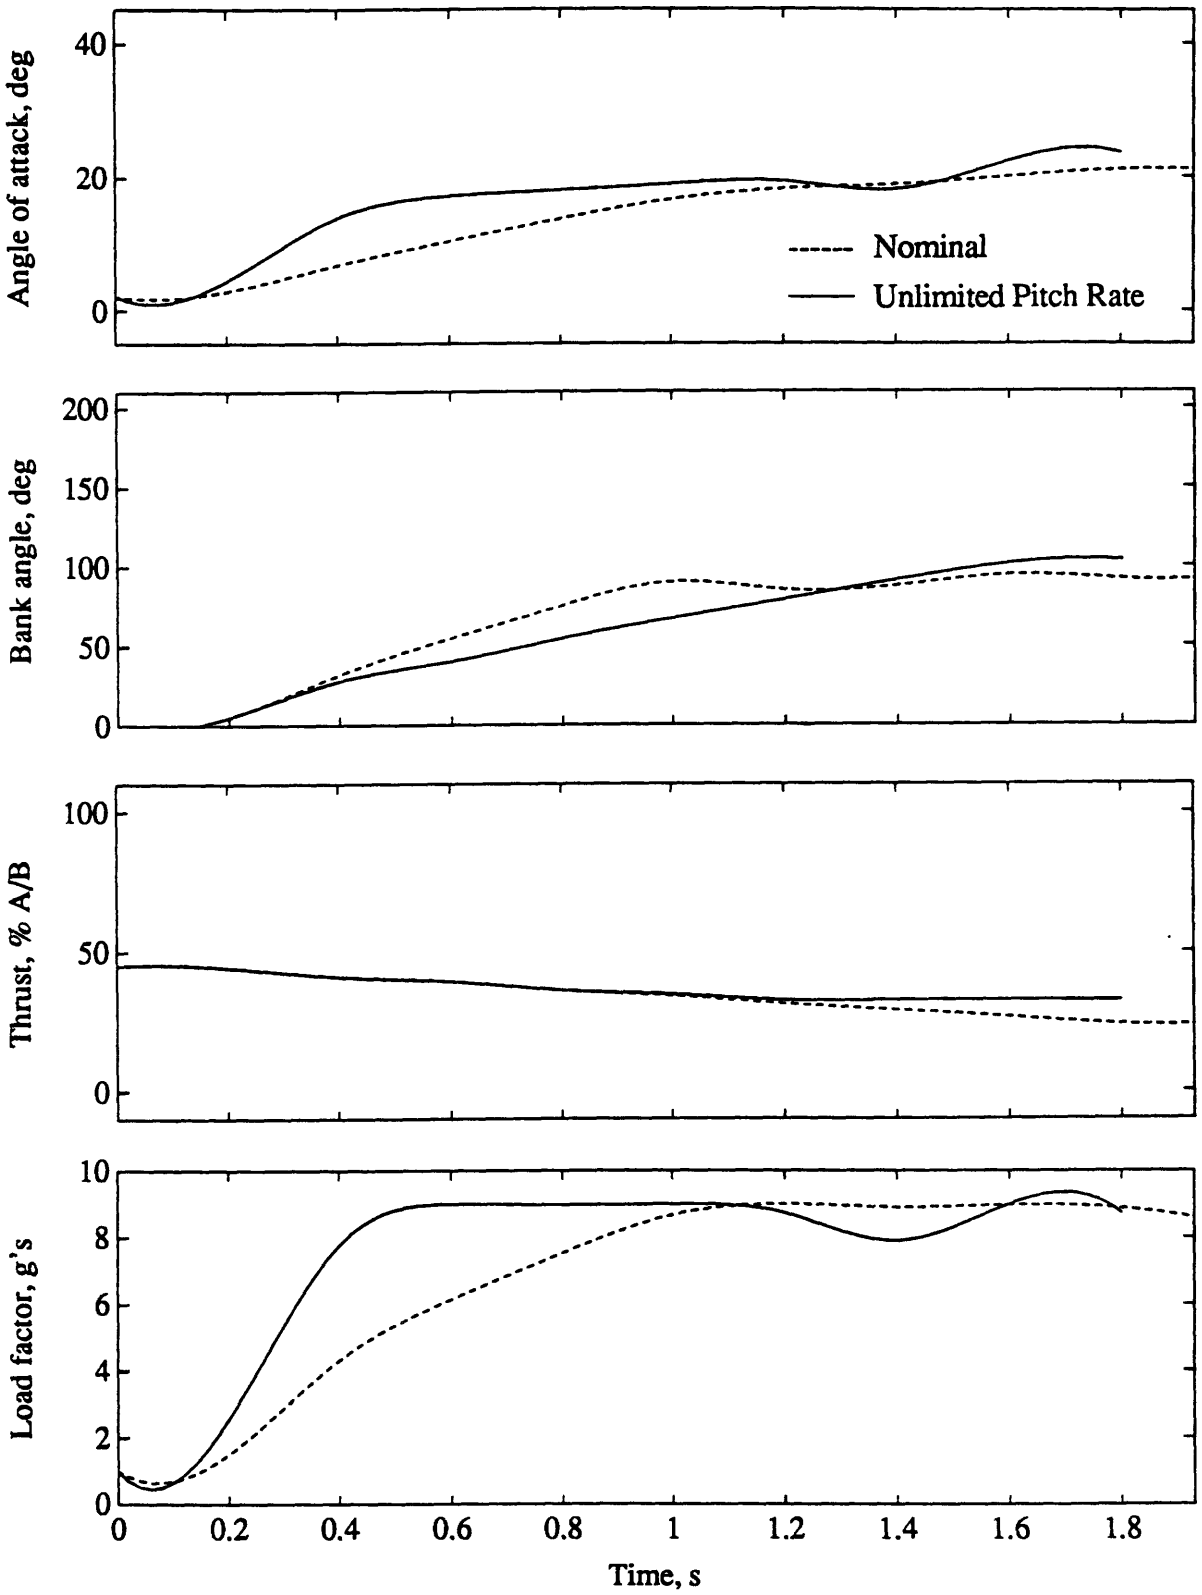

Figure 4.18. Optimal trajectories for nominal model vs. unlimited pitch rate model to (45,0), a short horizontal turn. The trajectory is so short that speed stays up and α_{max} is never reached to prevent load factor limit violations. Note the good correspondence between the α plateau and maximum load factor.

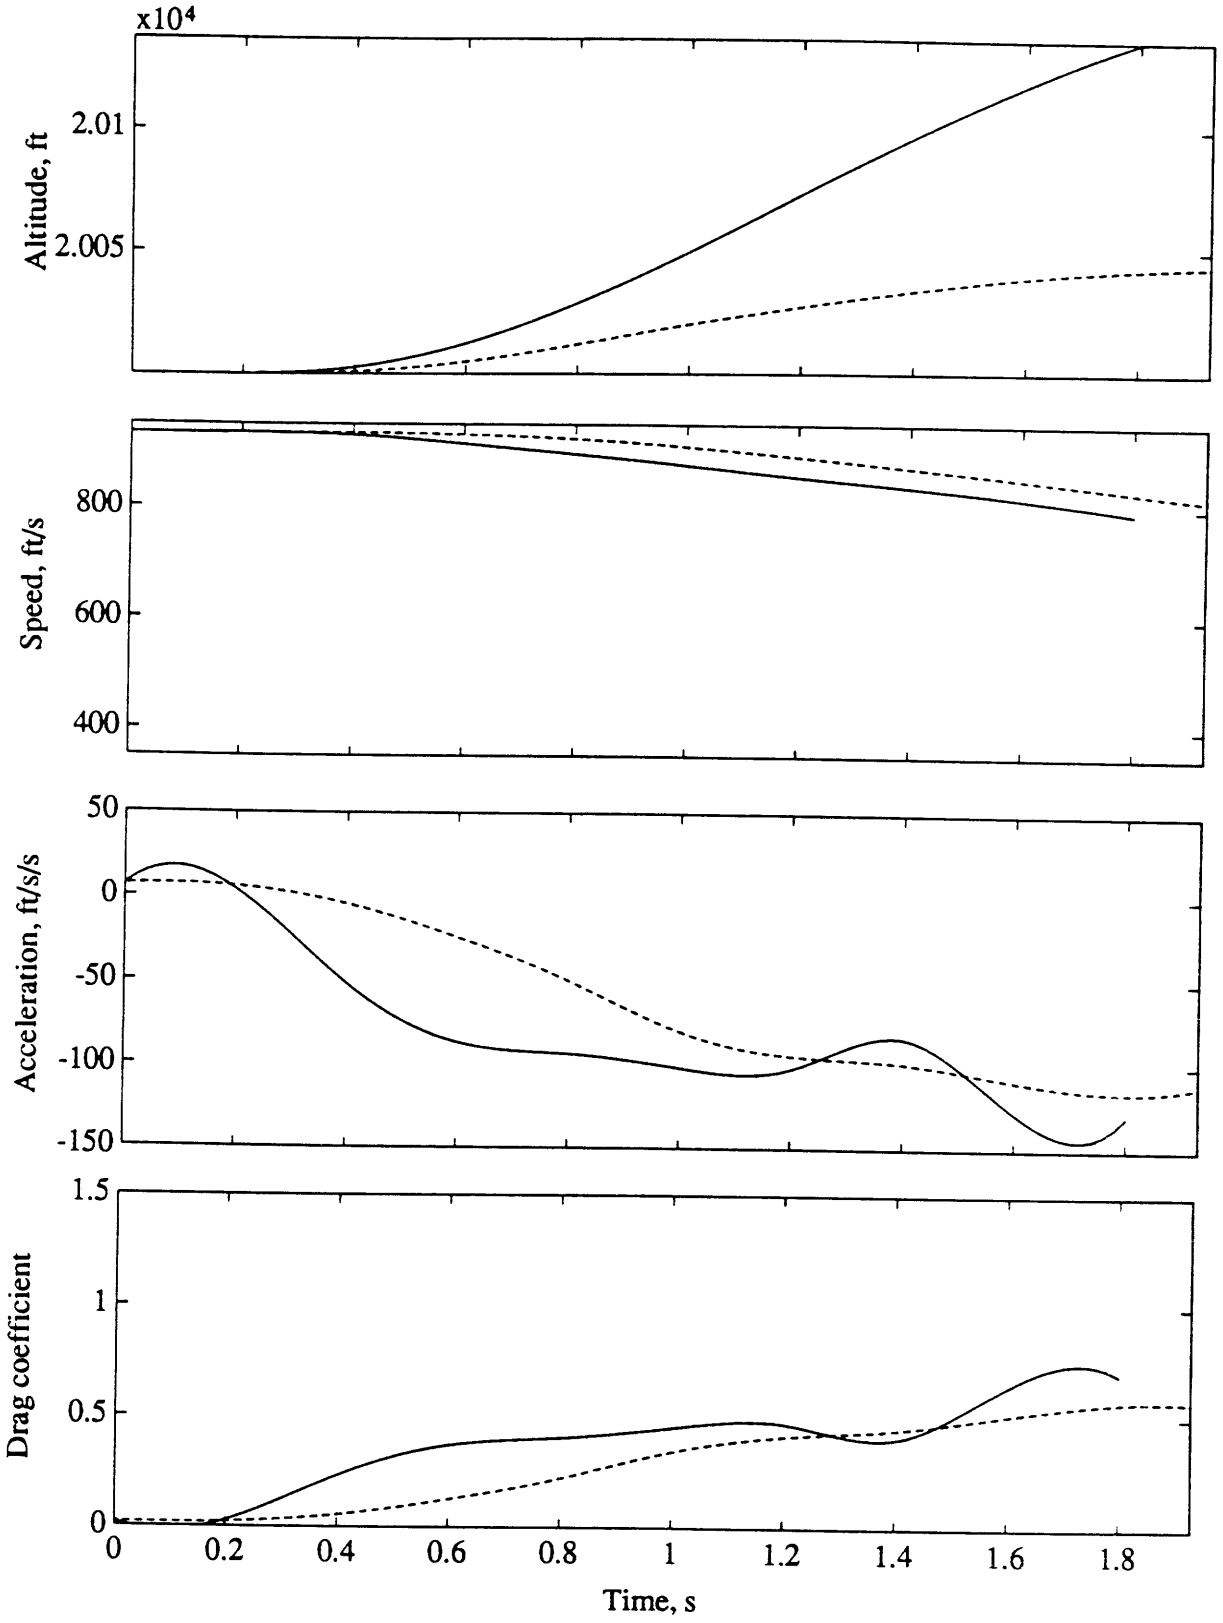

Figure 4.18 (cont'd)

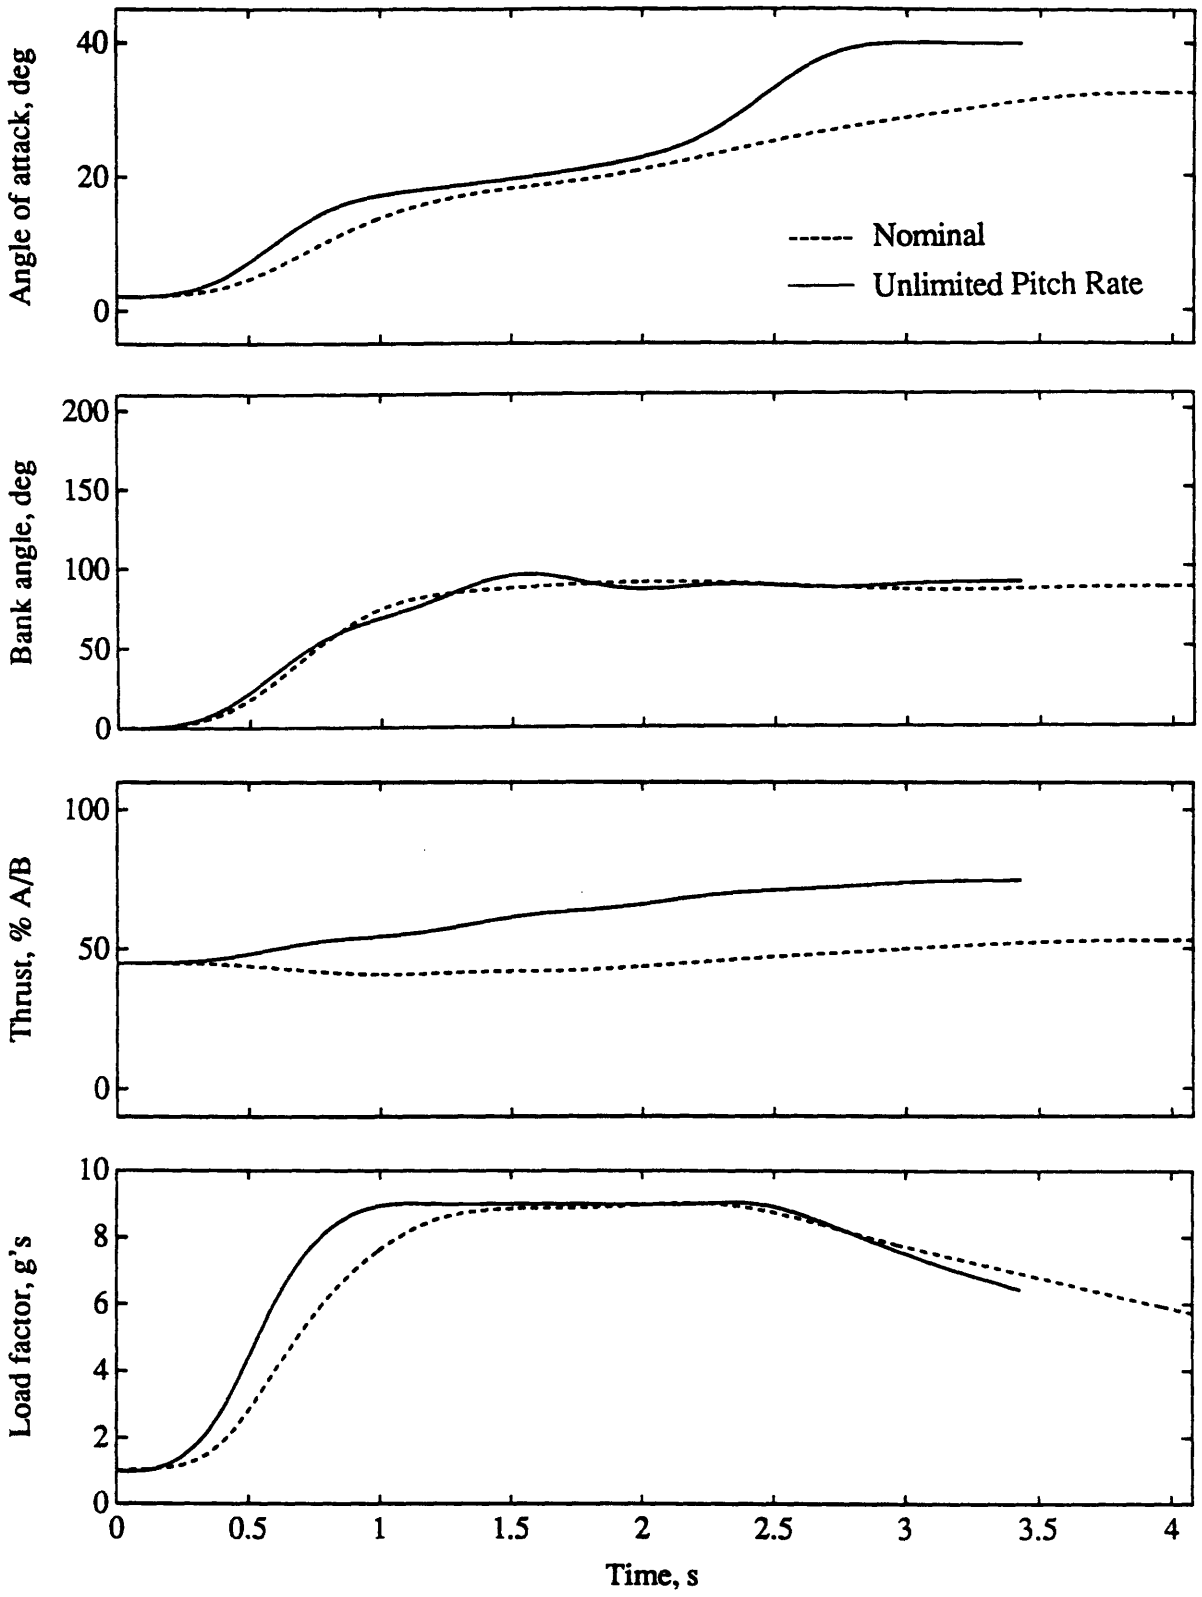

Figure 4.19. Optimal trajectories for nominal model vs. unlimited pitch rate model to (90,0), a moderate length horizontal turn. Angle of attack stays fairly low until the pop-up to α_{max} at the end where there is no time for drag to affect speed adversely.

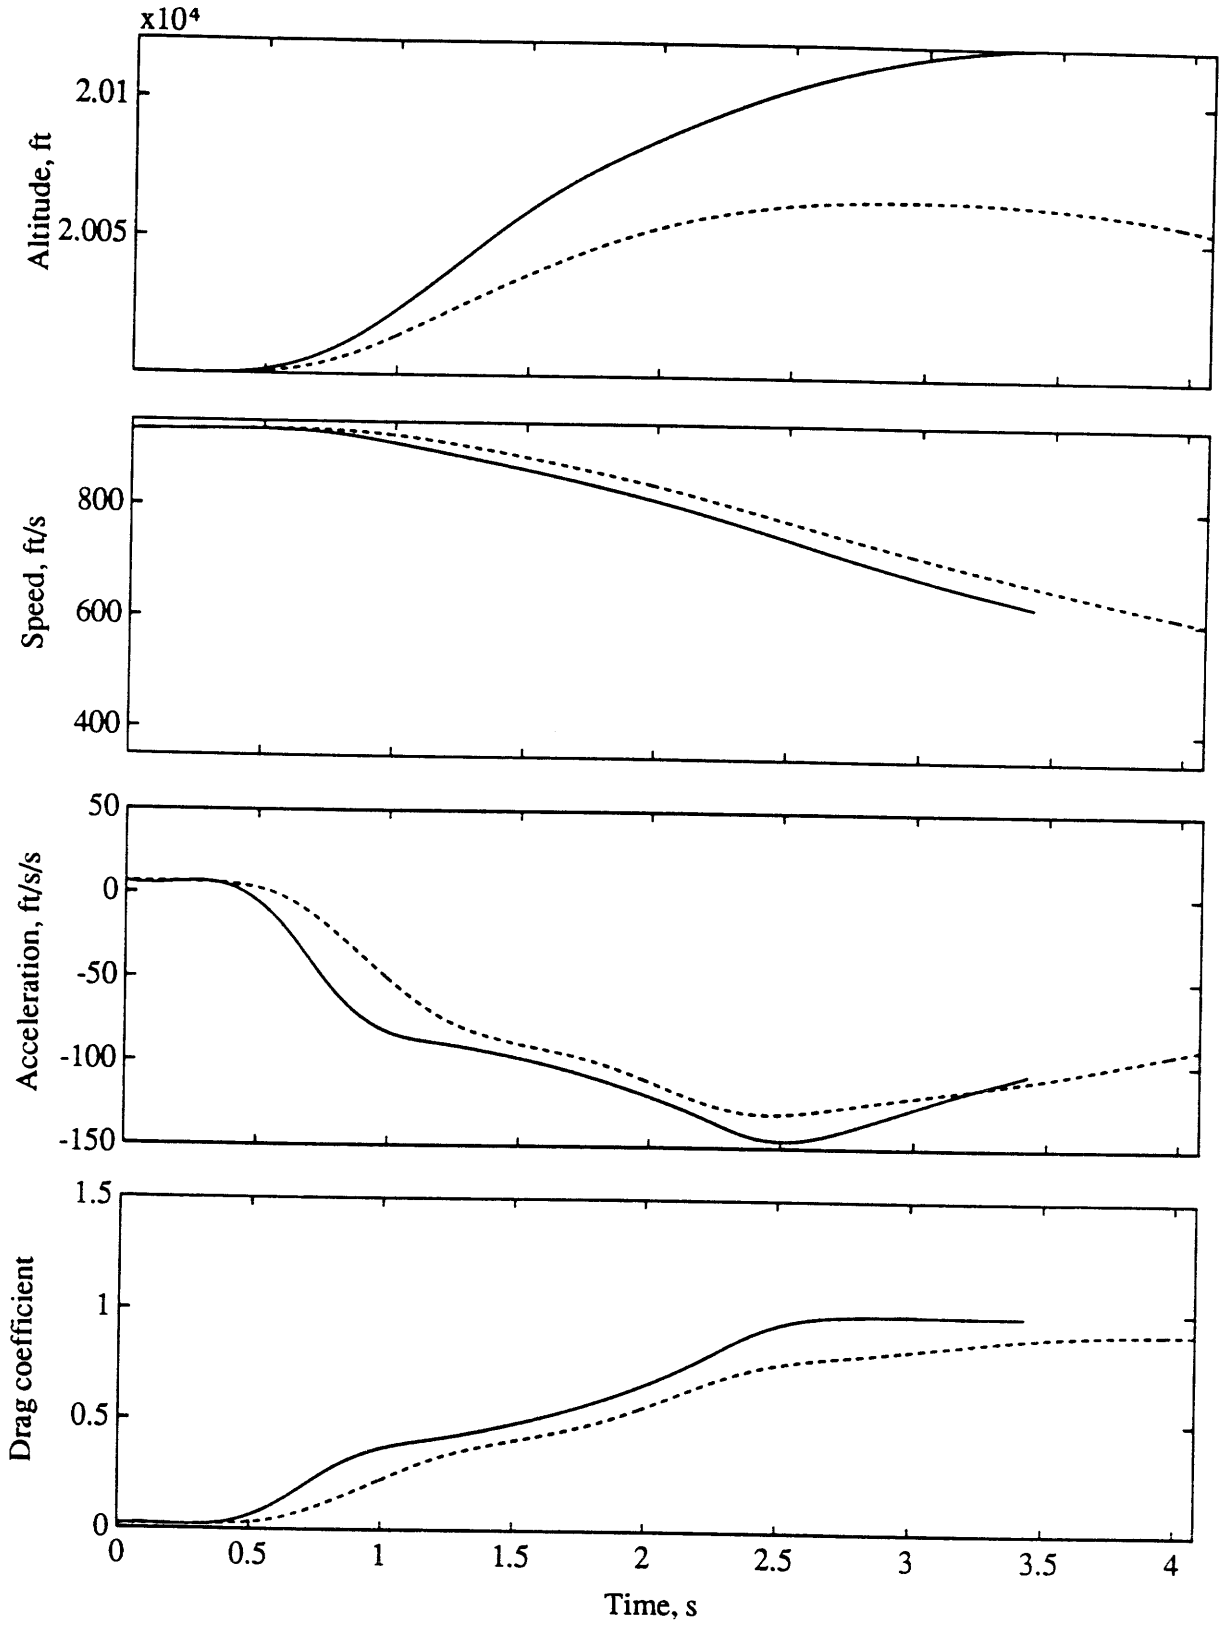

Figure 4.19 (cont'd)

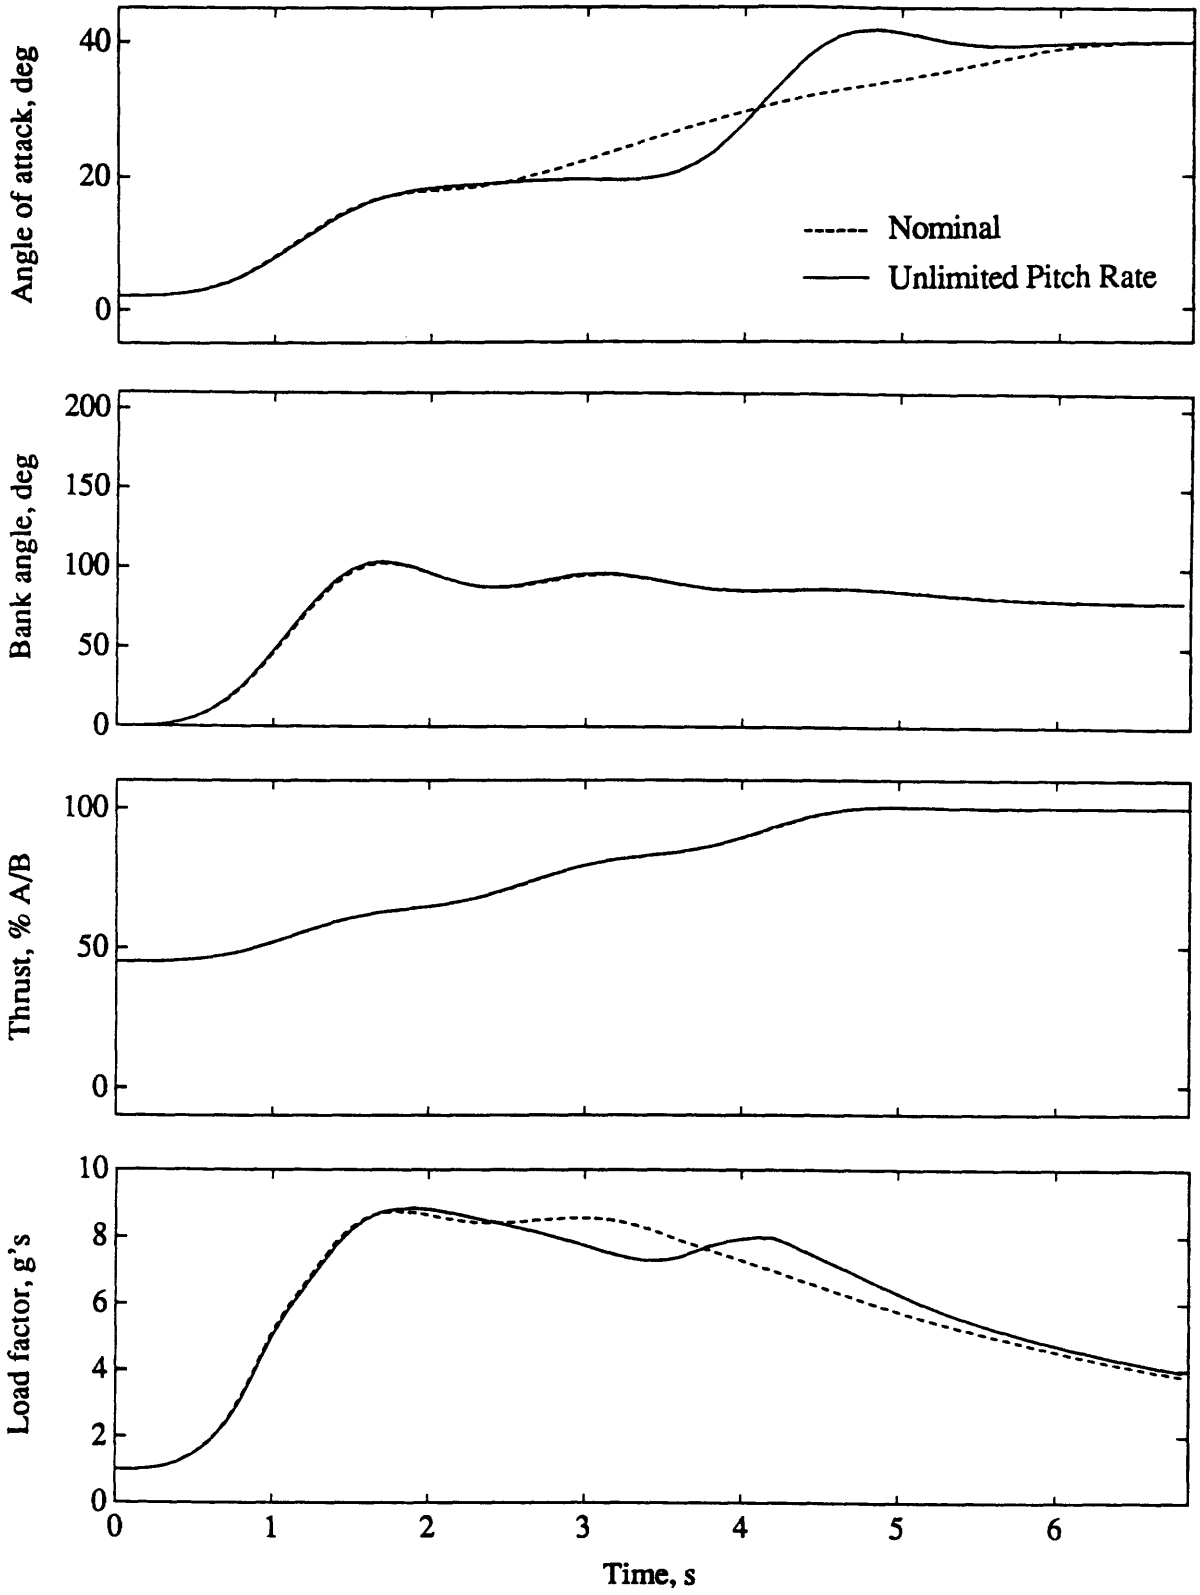

Figure 4.20. Optimal trajectories for nominal model vs. unlimited pitch rate model to (135,0), a long horizontal turn. The α plateau and maximum load factor correspondence is not good this time. The angle of attack stays low because of the drag divergence at the flight condition (Figure 3.2).

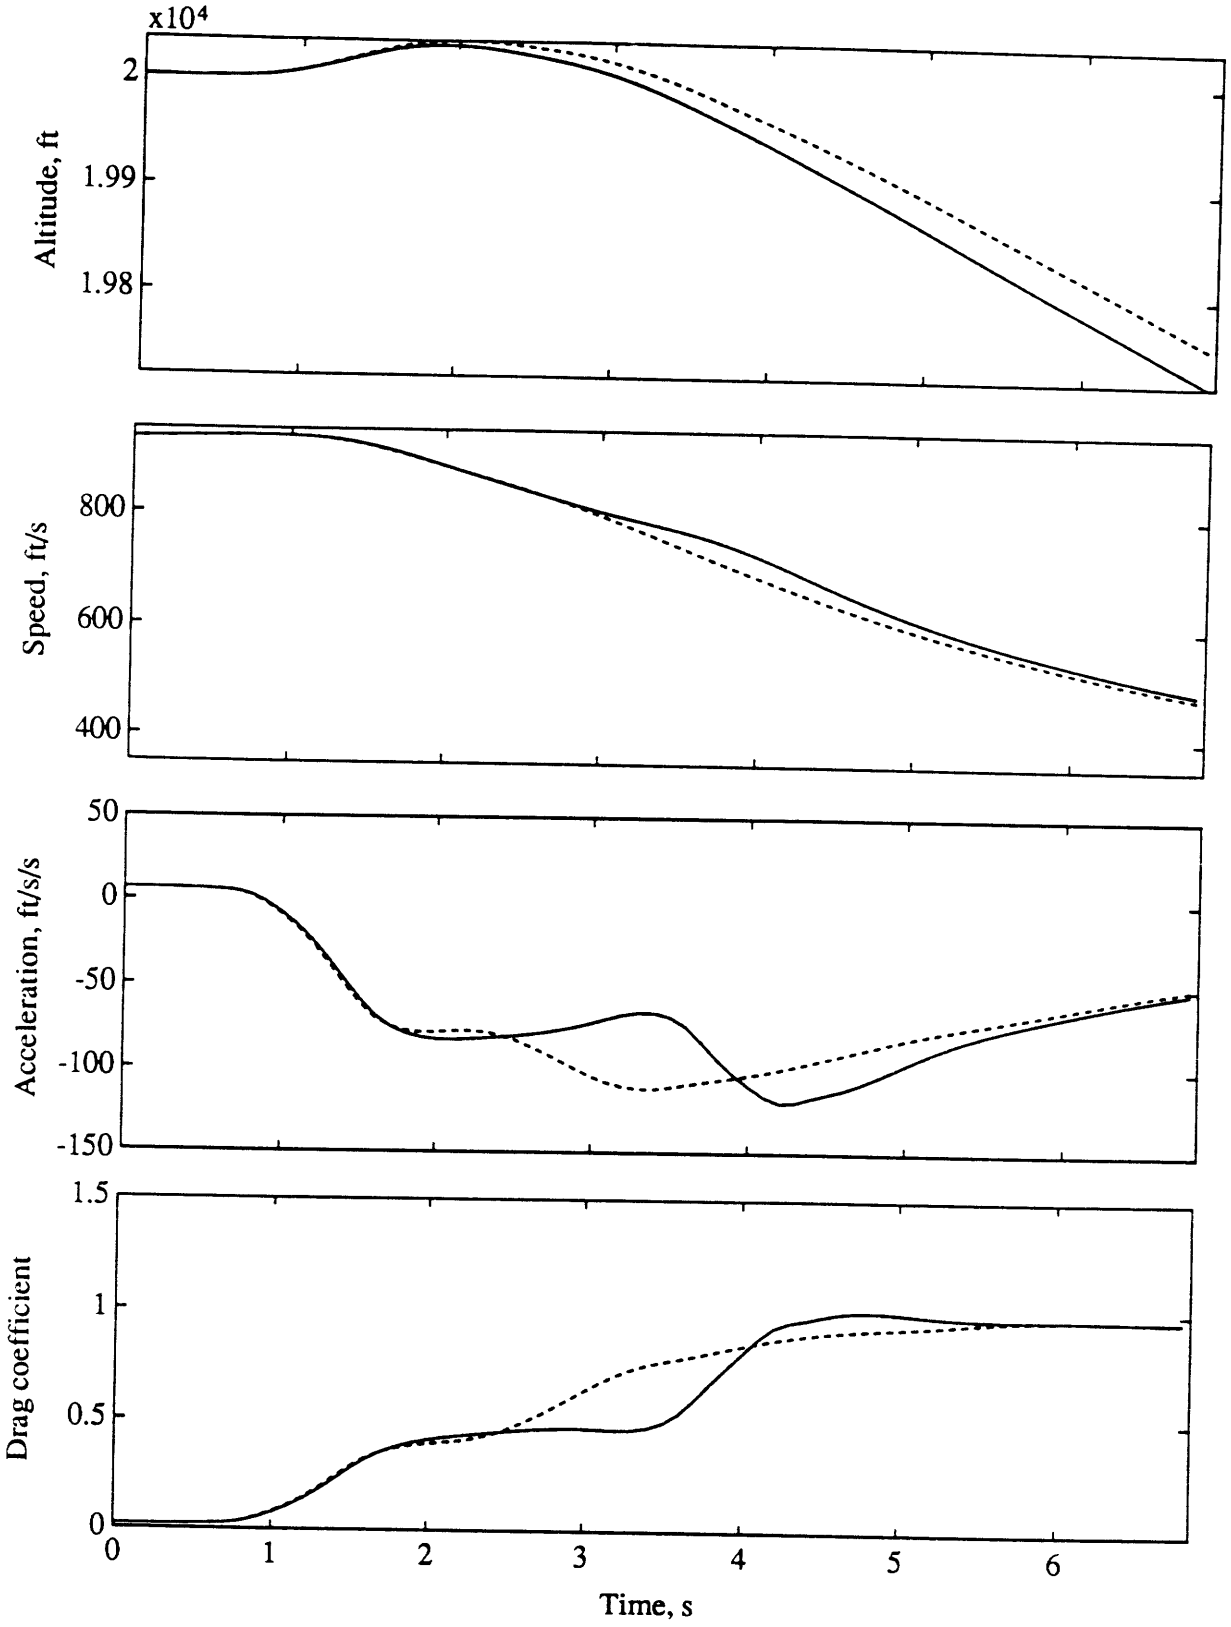

Figure 4.20 (cont'd)

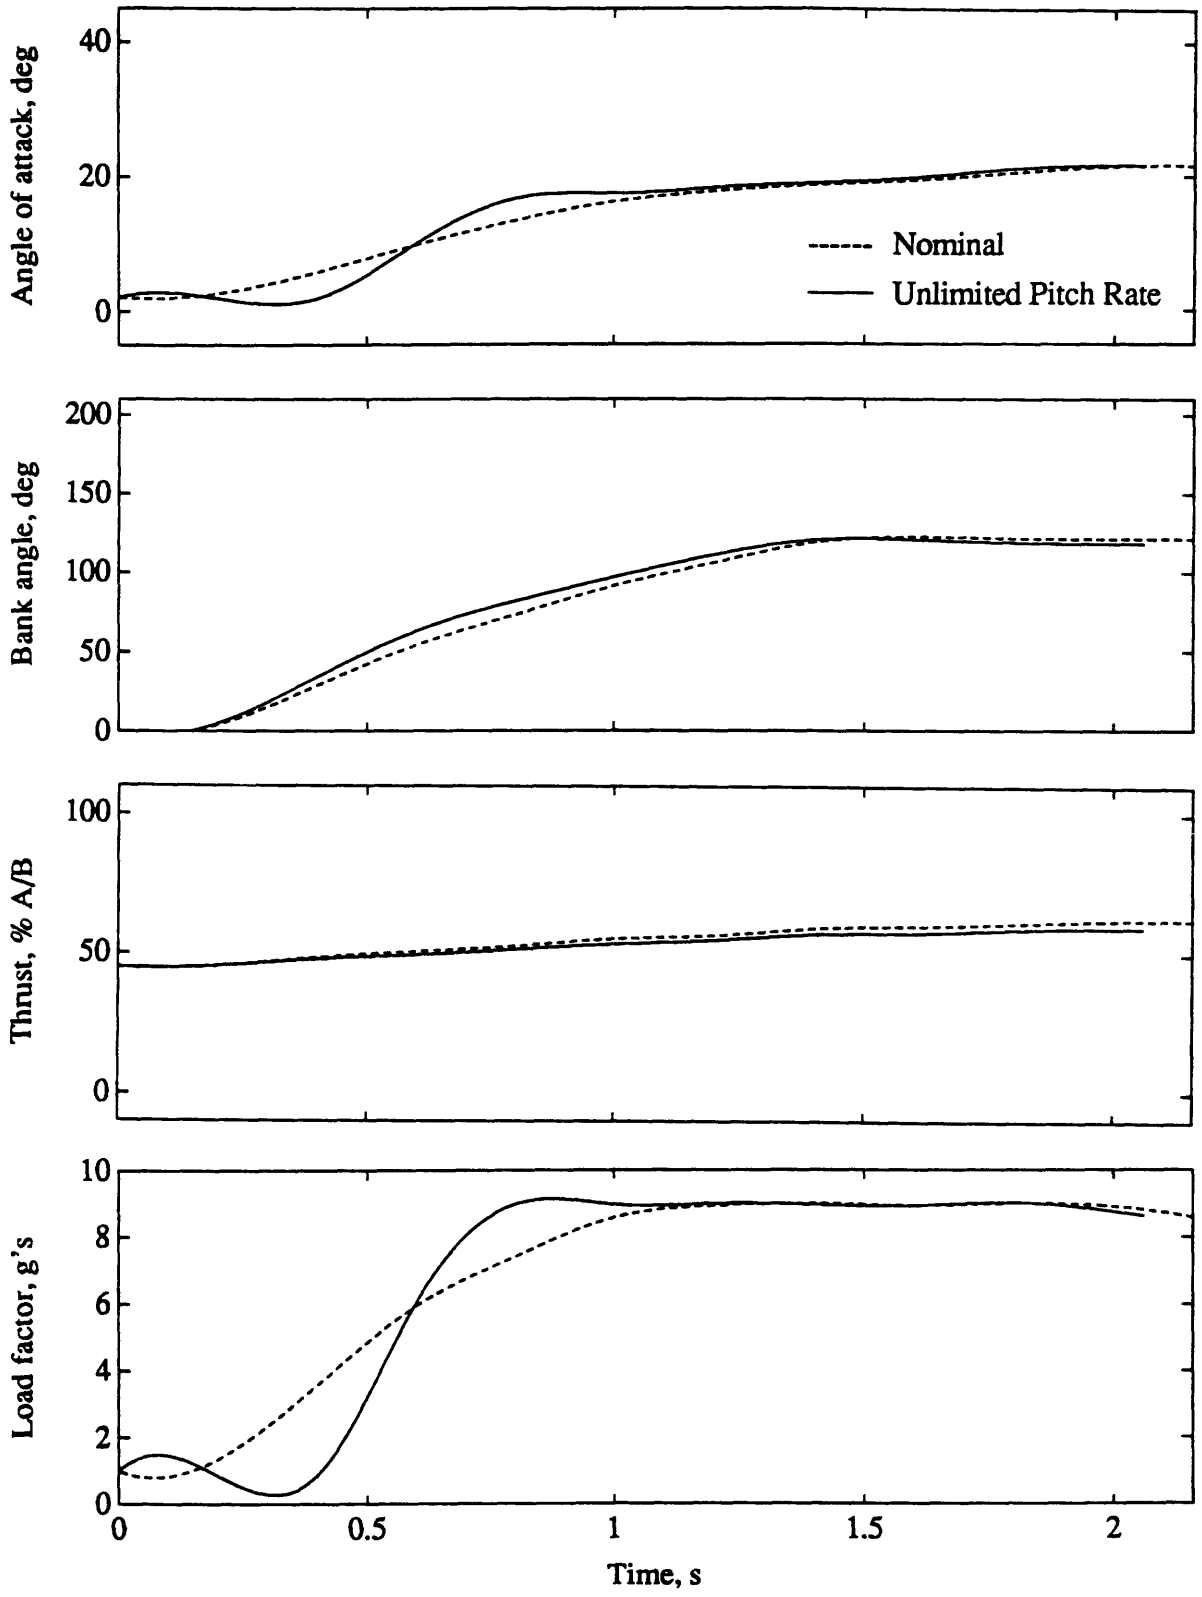

Figure 4.21. Optimal trajectories for nominal model vs. unlimited pitch rate model to (45,-20), a short lower hemisphere trajectory. The unlimited model pitches after the nominal model because the target is below the aircraft. Any pitch before the roll is completed increases the elevation angle.

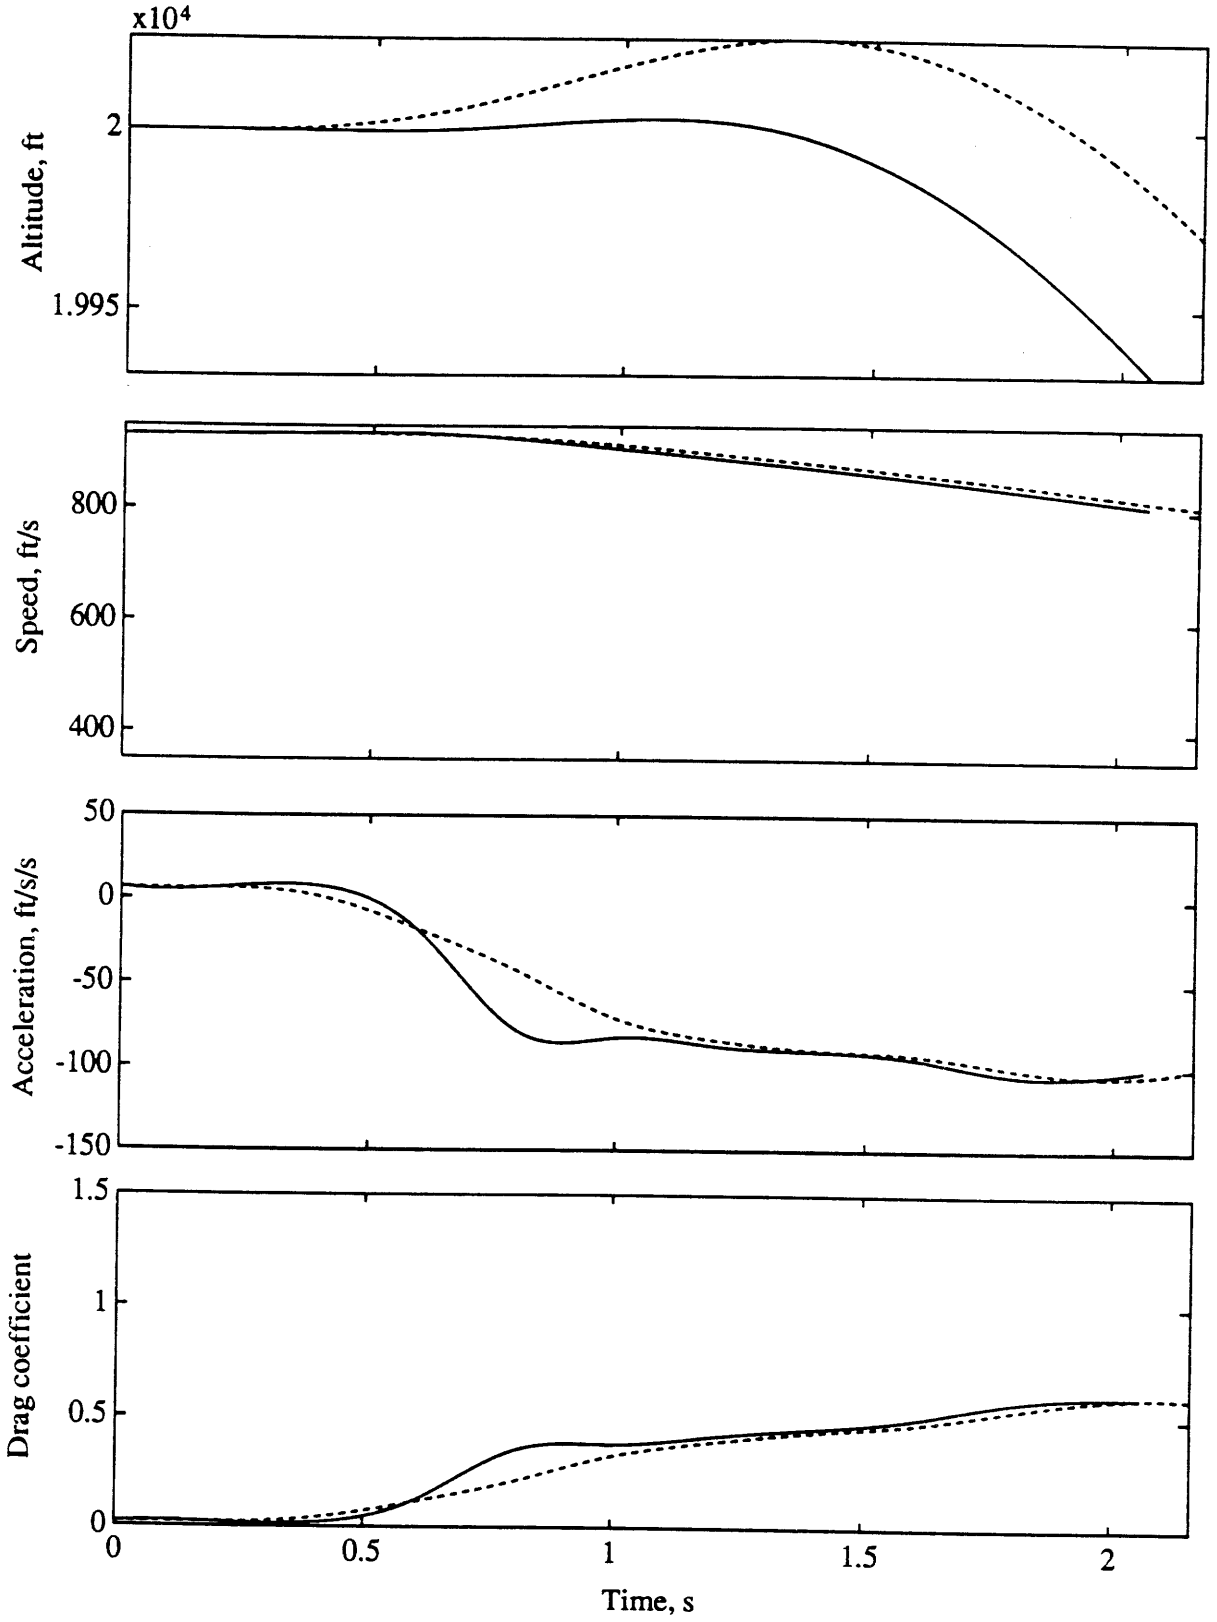

Figure 4.21 (cont'd)

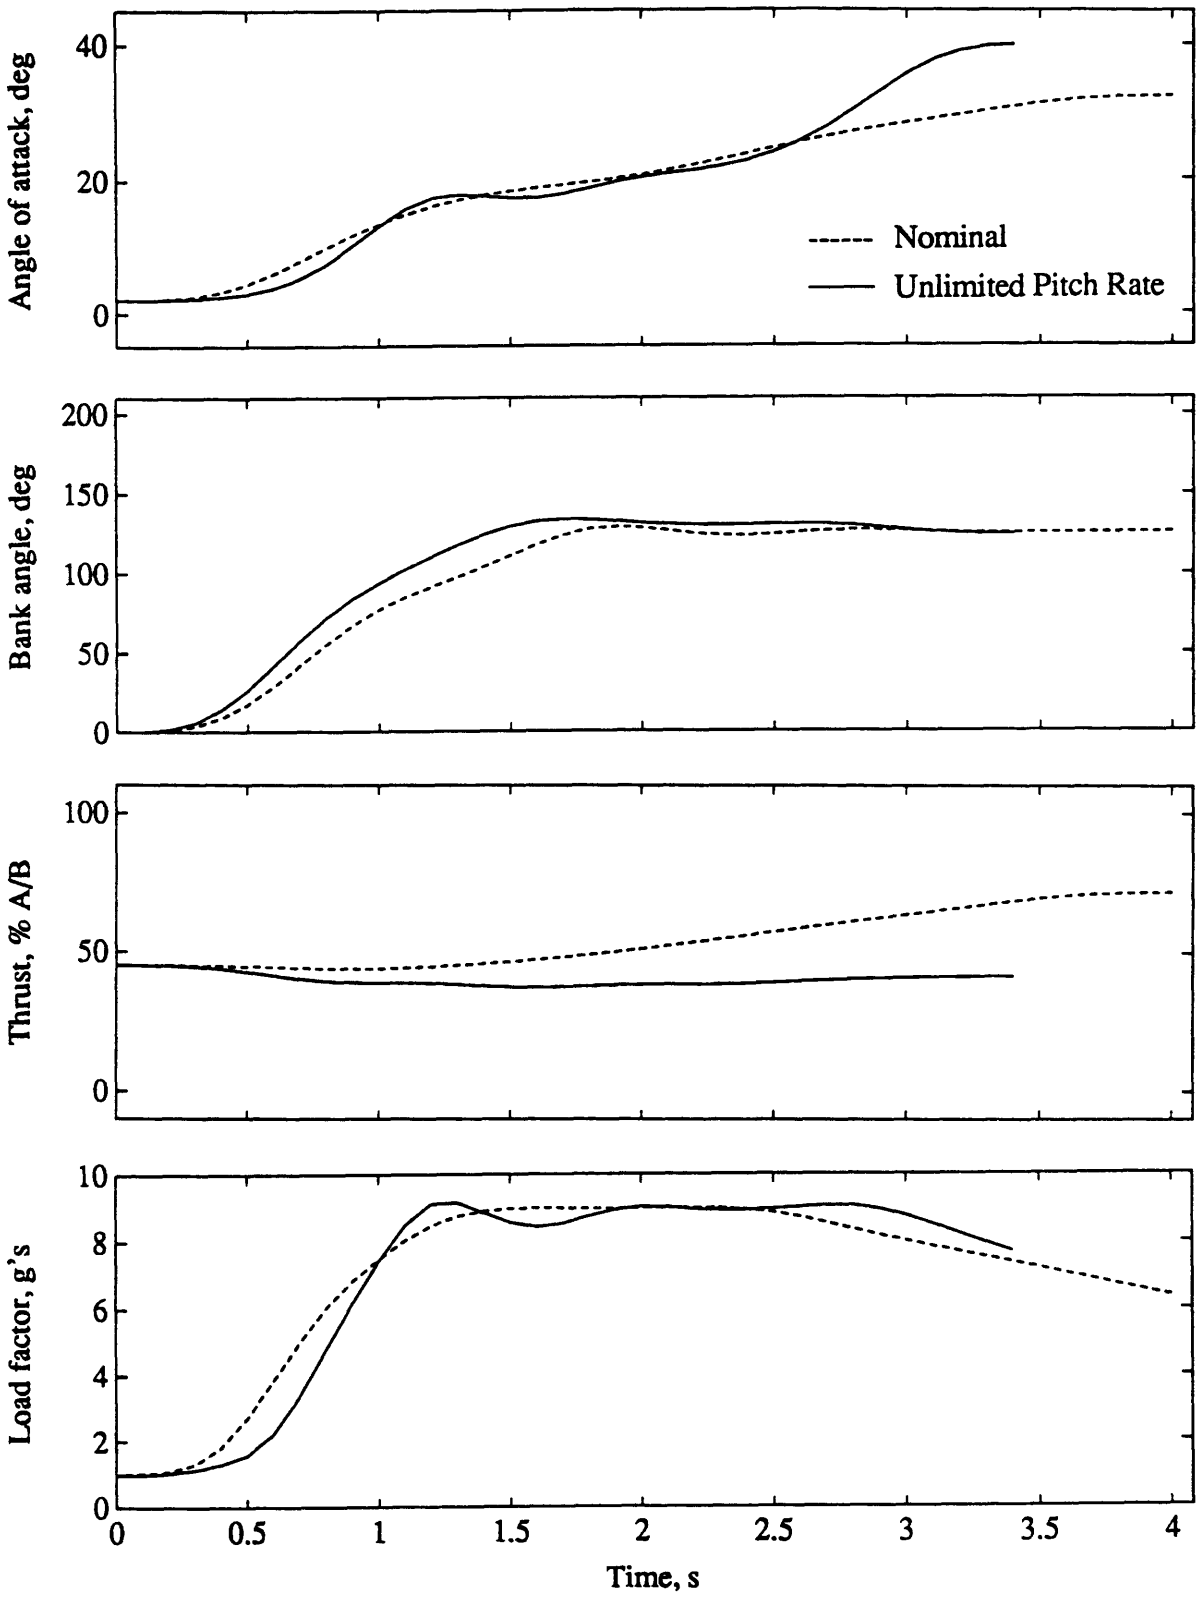

Figure 4.22. Optimal trajectories for nominal model vs. unlimited pitch rate model to (90,-45), a moderate length lower hemisphere trajectory. The α pop-up is present because there is no danger of violating the load factor limit.

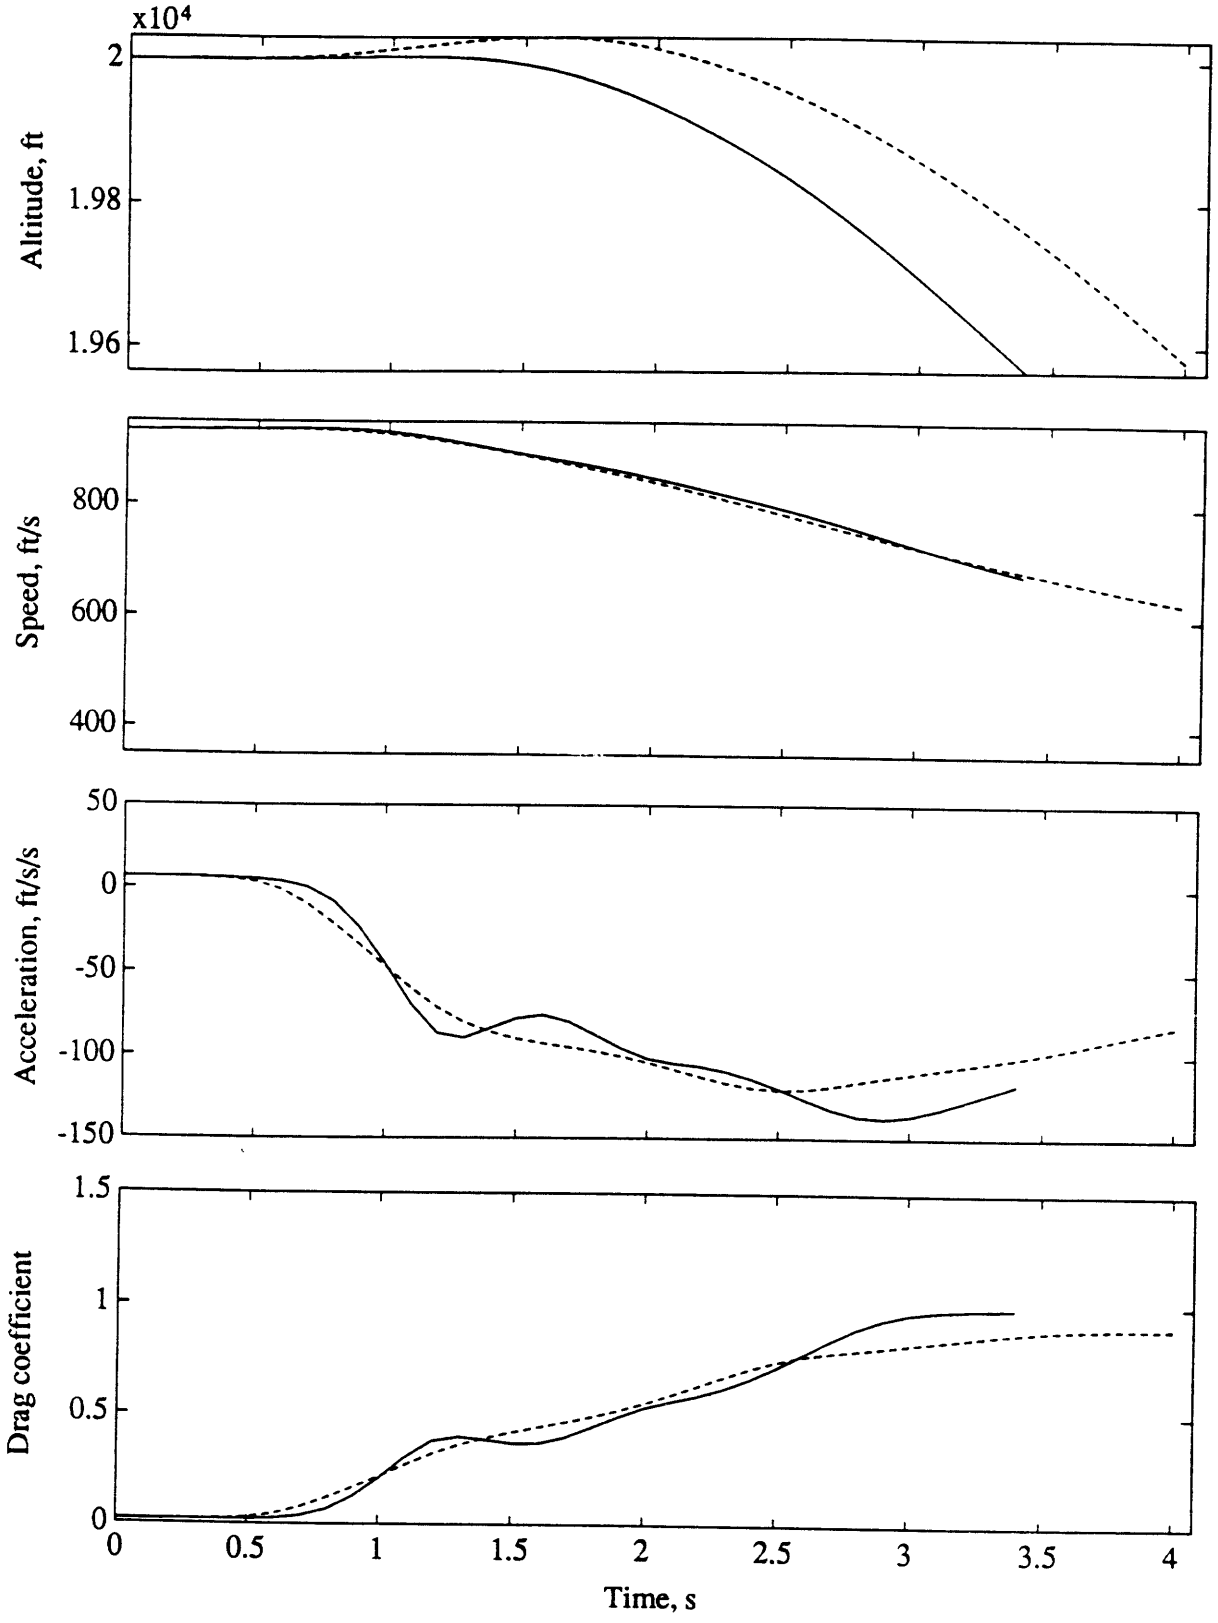

Figure 4.22 (cont'd)

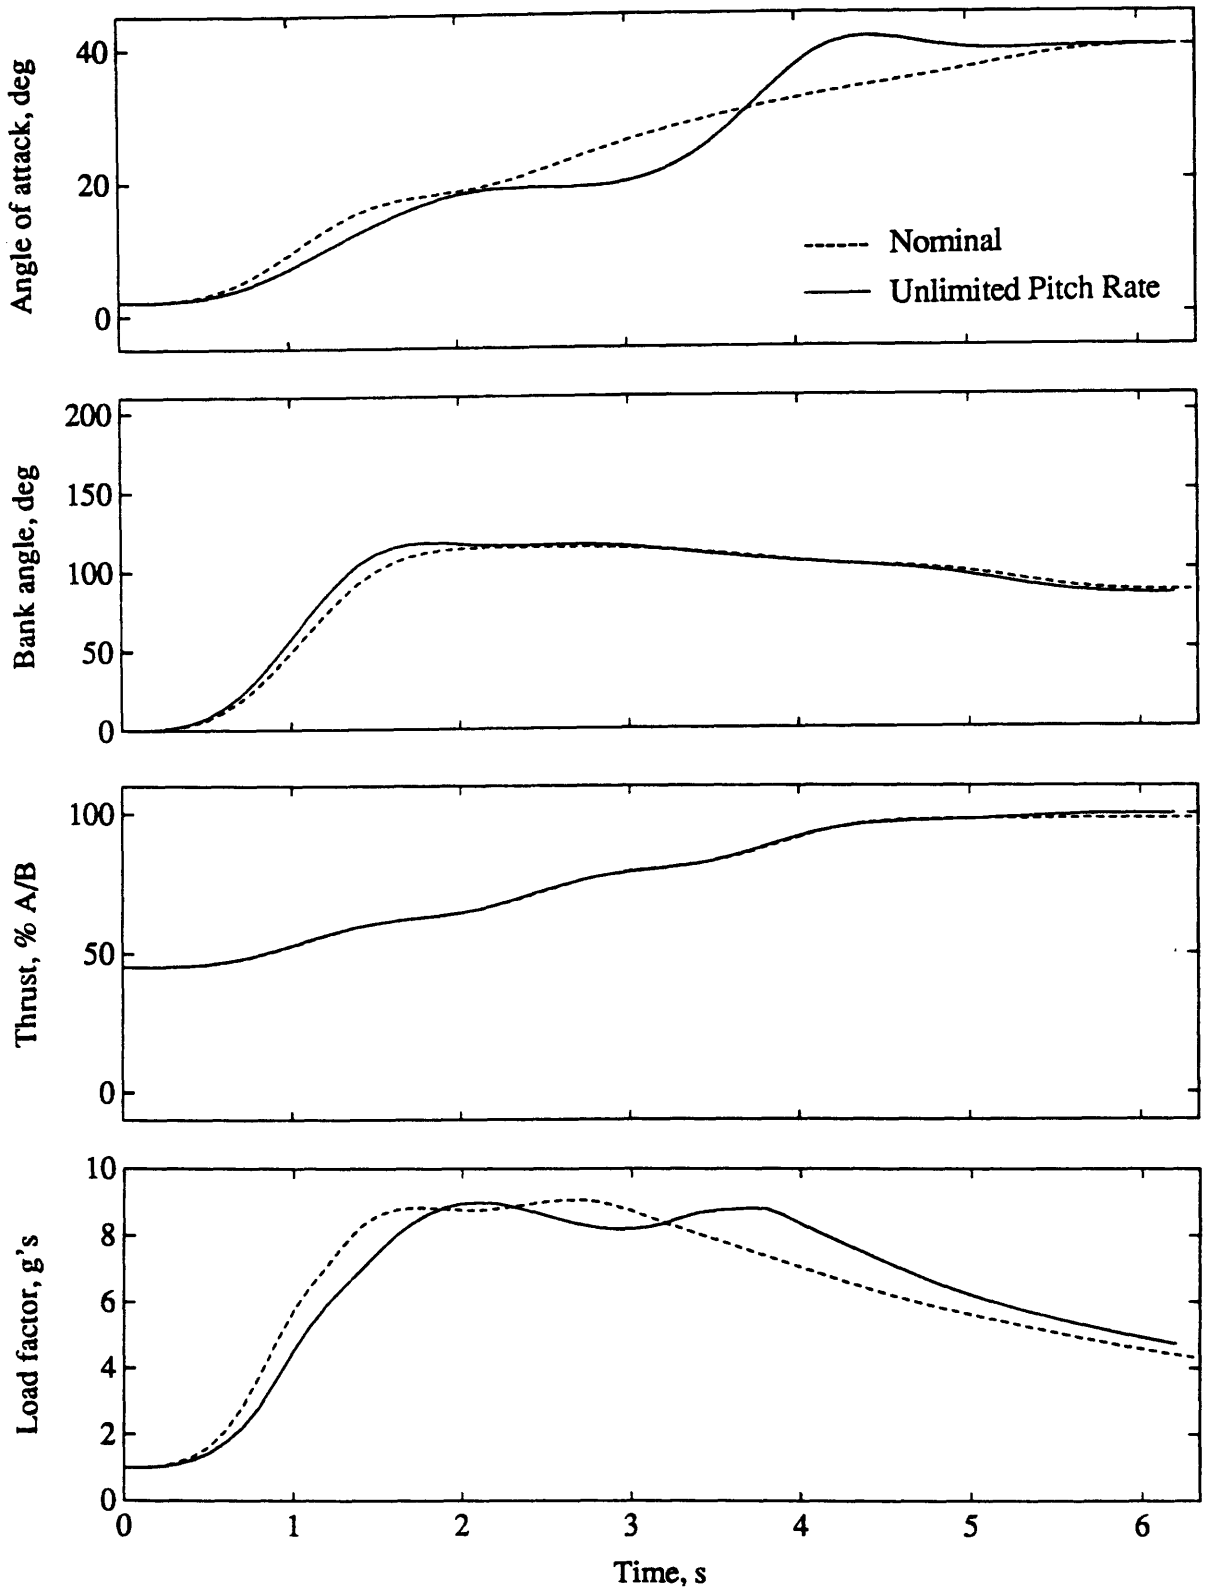

Figure 4.23. Optimal trajectories for nominal model vs. unlimited pitch rate model to (135,-20), a long lower hemisphere trajectory. Energy management is important for long trajectories, so the unlimited α and drag are lower than for the nominal case during most of the trajectory.

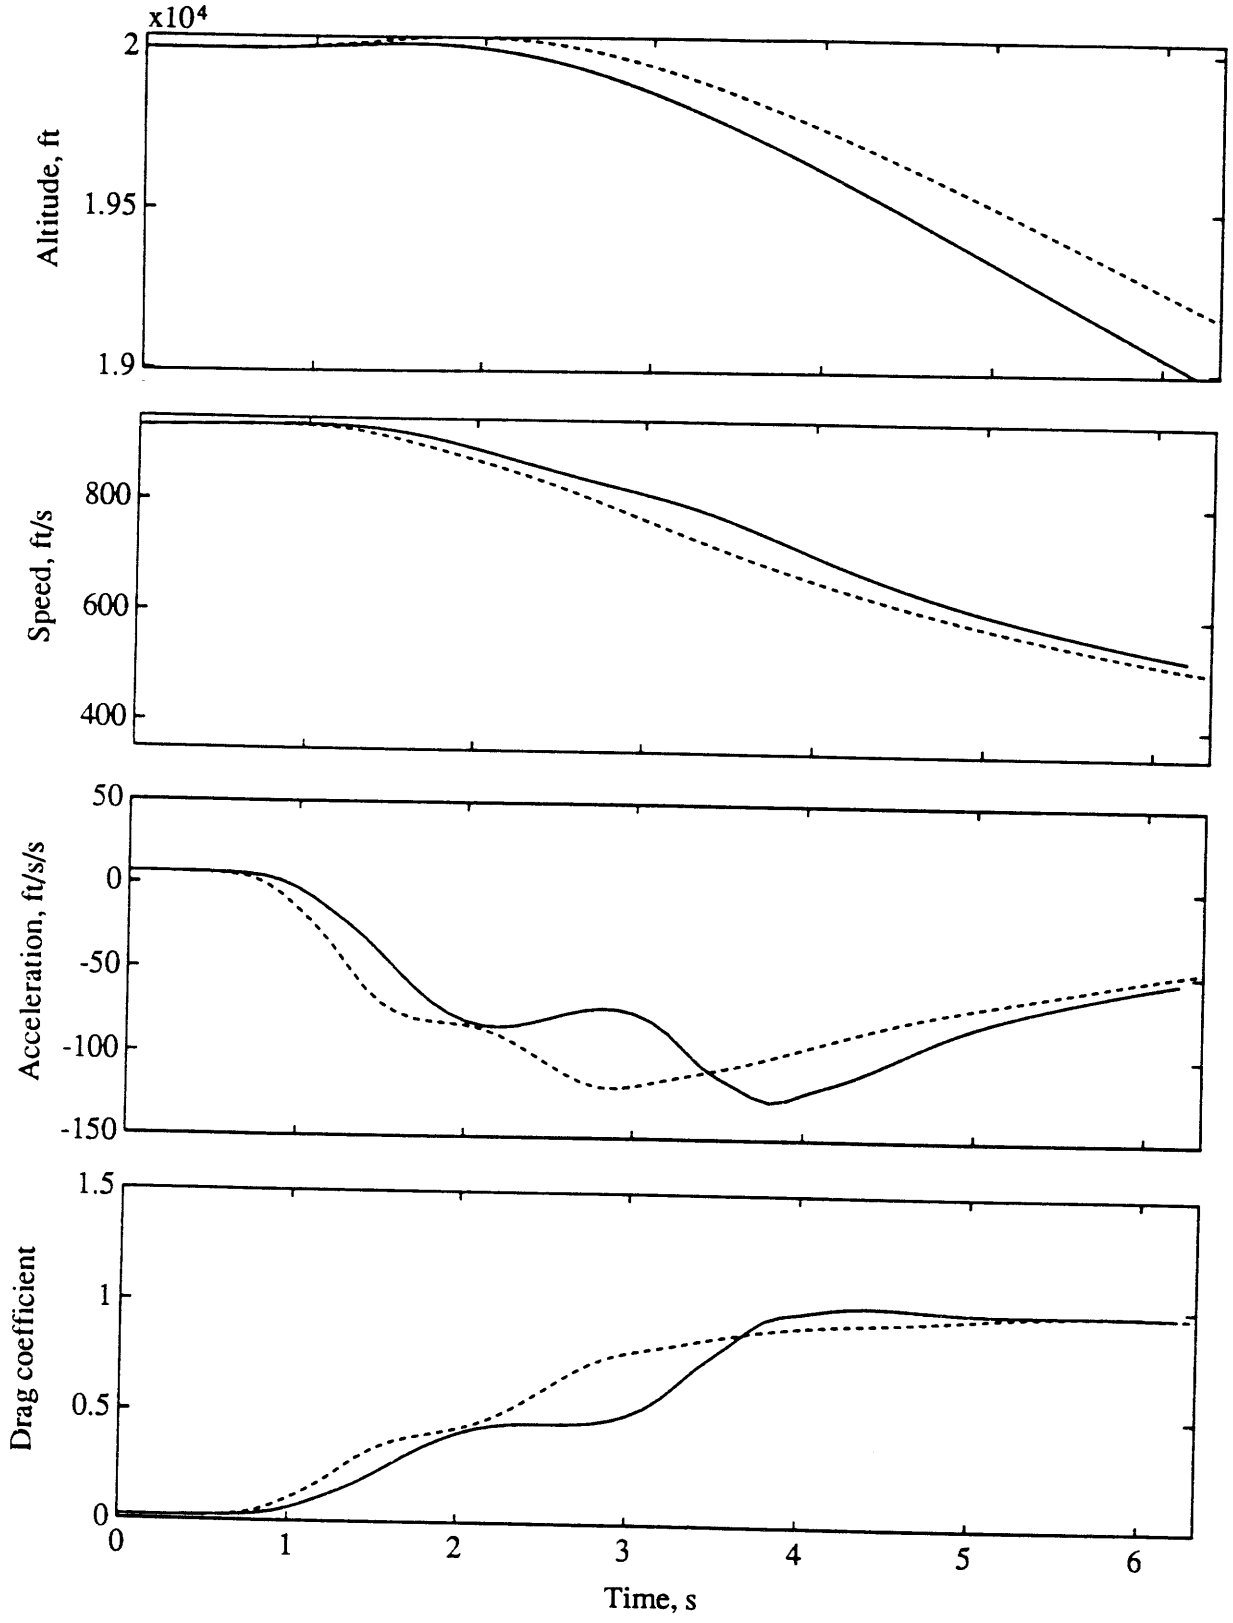

Figure 4.23 (cont'd)

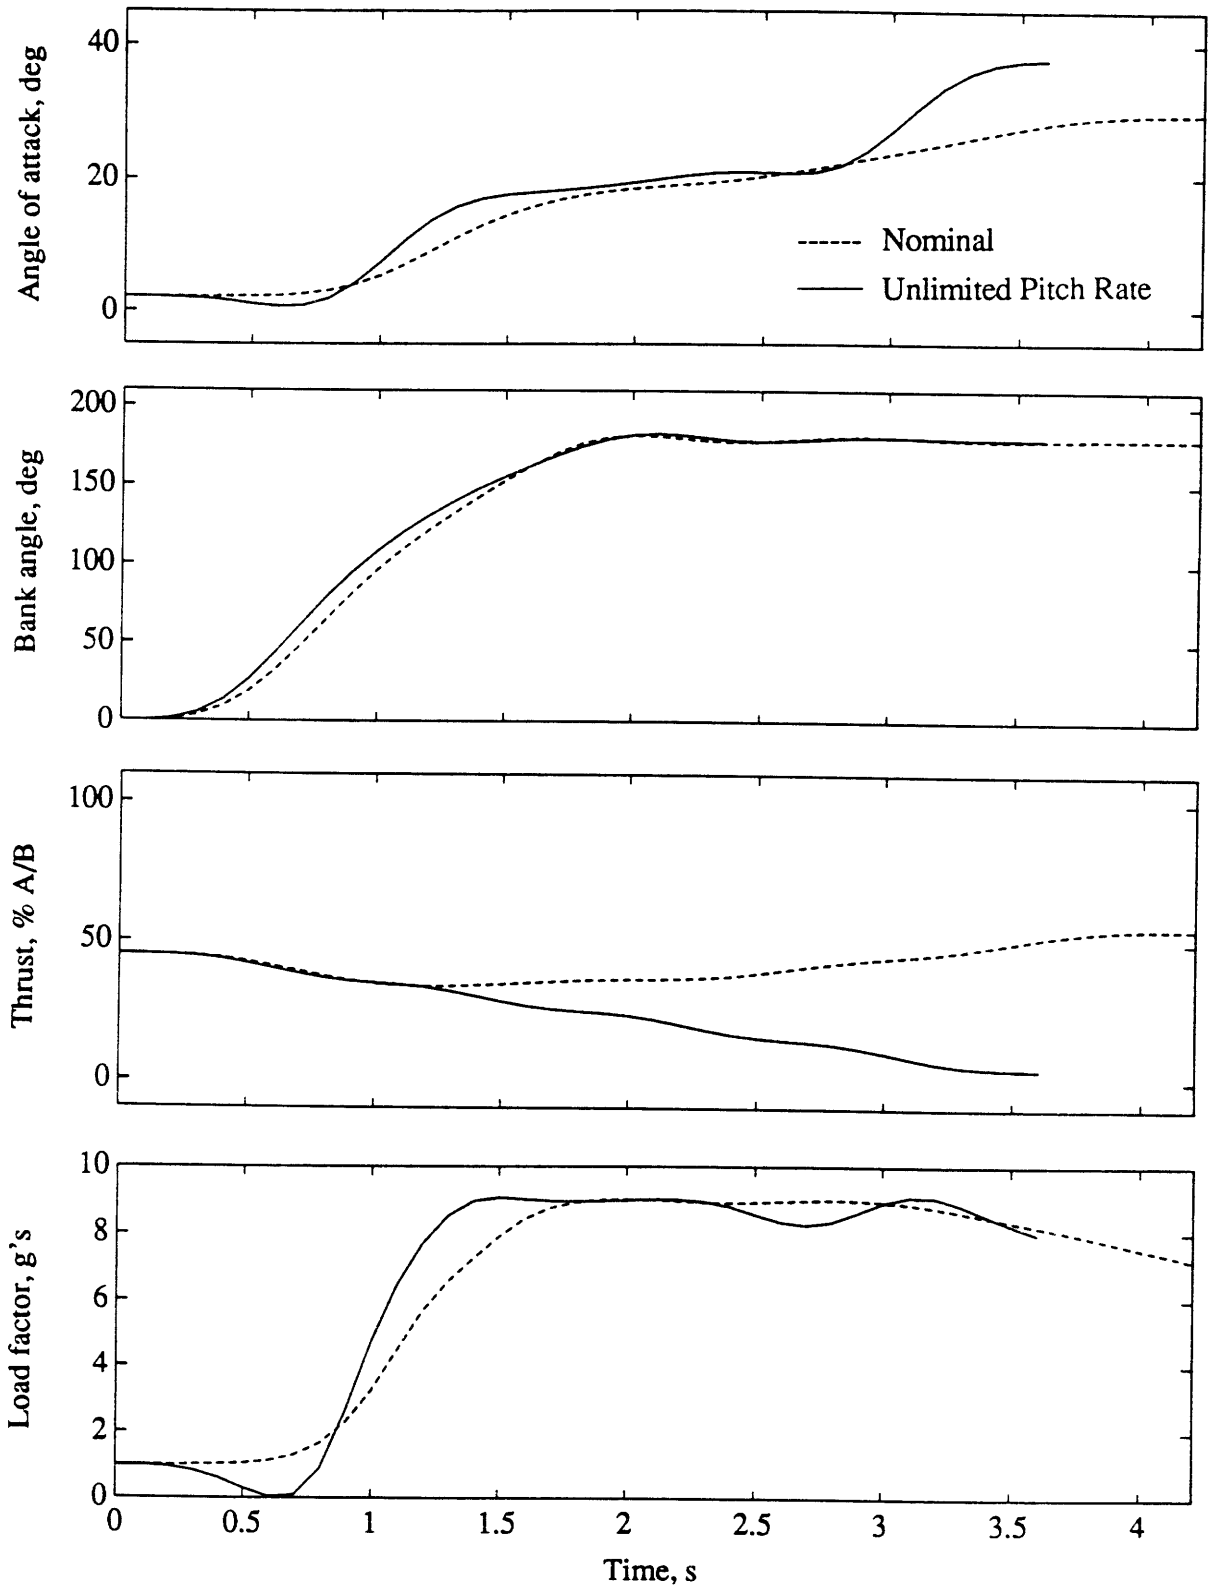

Figure 4.24. Optimal trajectories for nominal model vs. unlimited pitch rate model to (15,-88), nearly directly beneath the aircraft. Unlimited thrust reaches idle to allow the angle of attack pop-up.

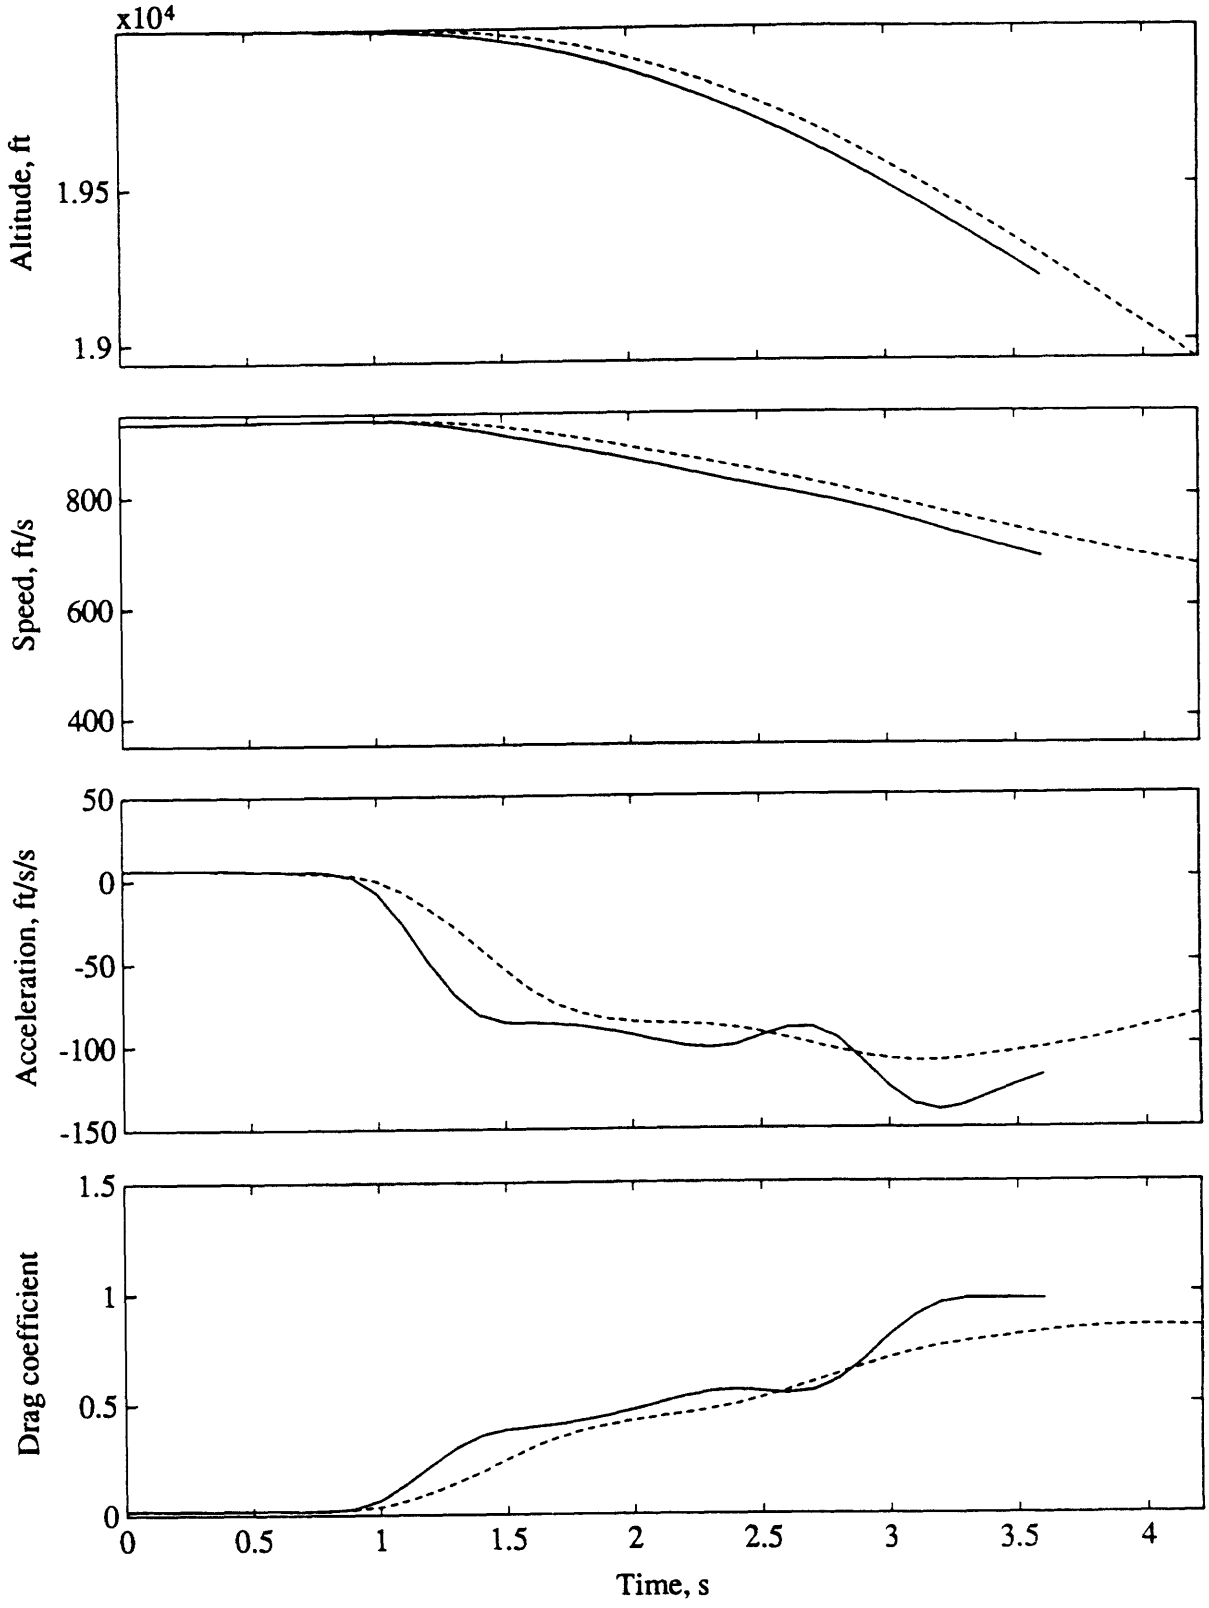

Figure 4.24 (cont'd)

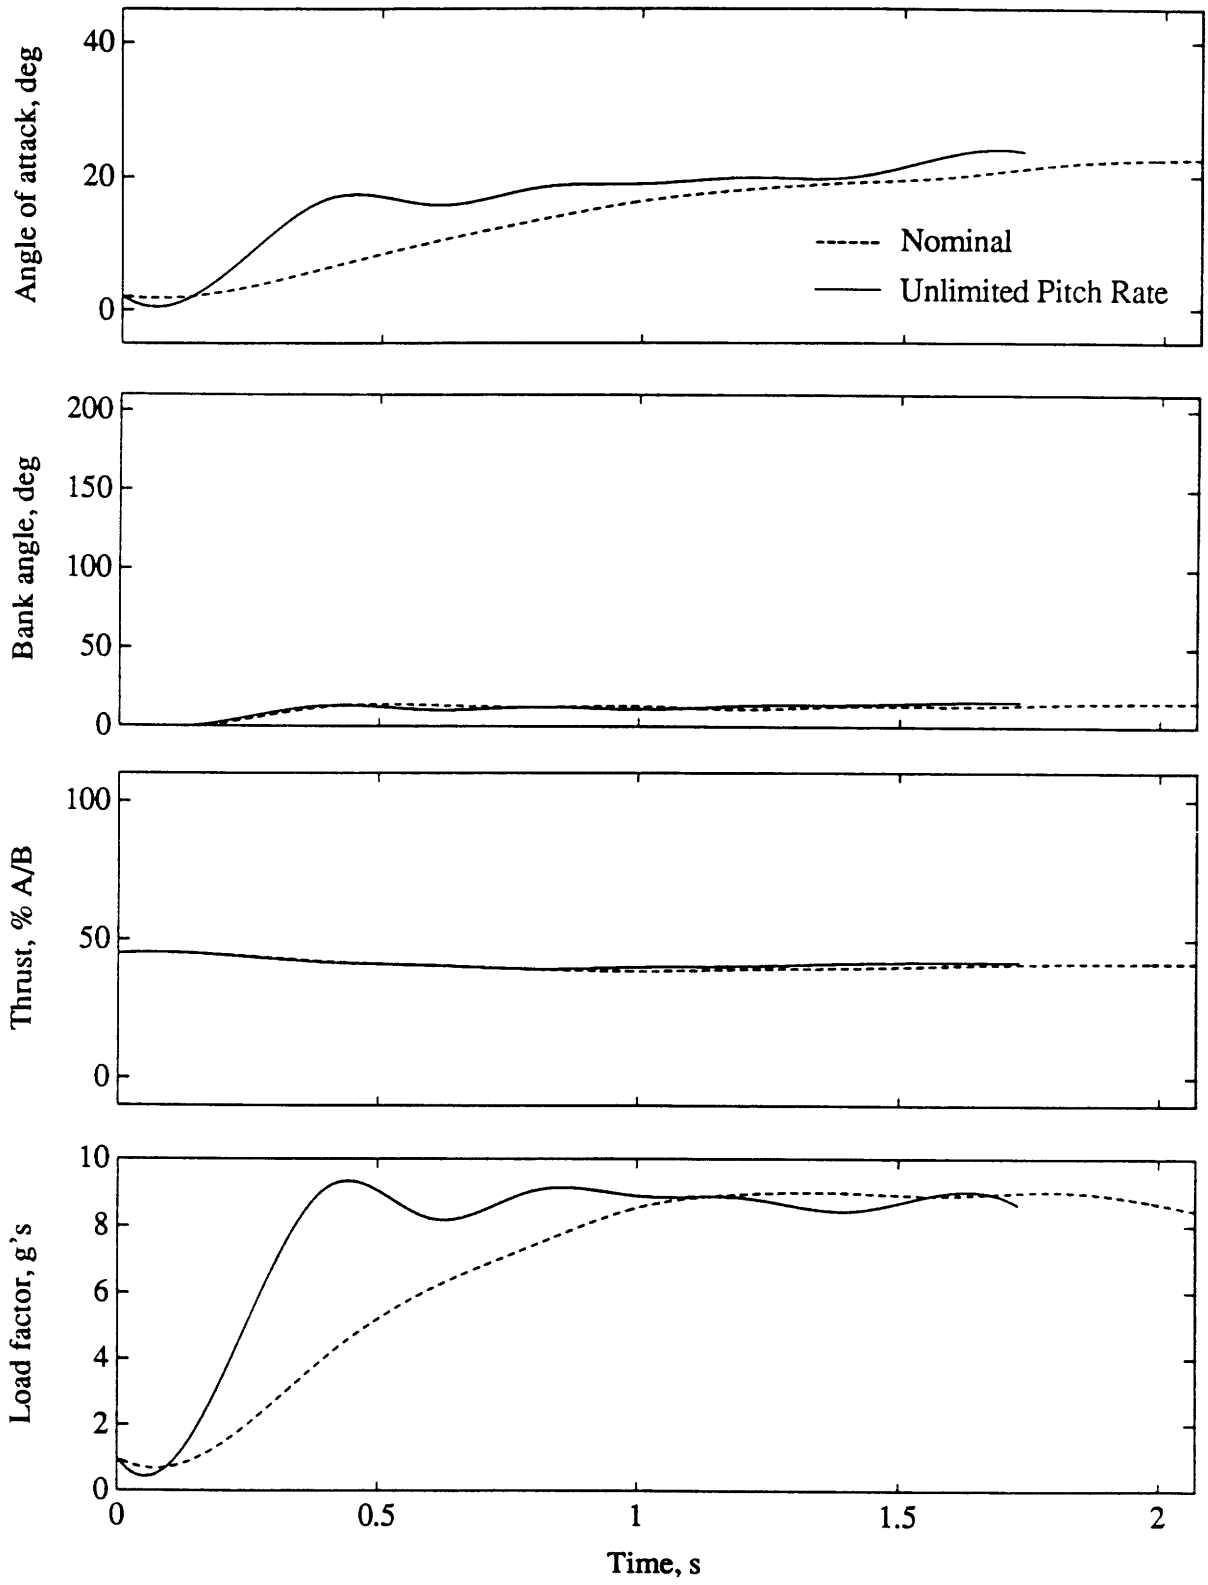

Figure 4.25. Optimal trajectories for nominal model vs. unlimited pitch rate model to (15,45), a forward upper hemisphere target. The initial α is much higher for the unlimited model because the trajectory is so short that there are no great speed penalties.

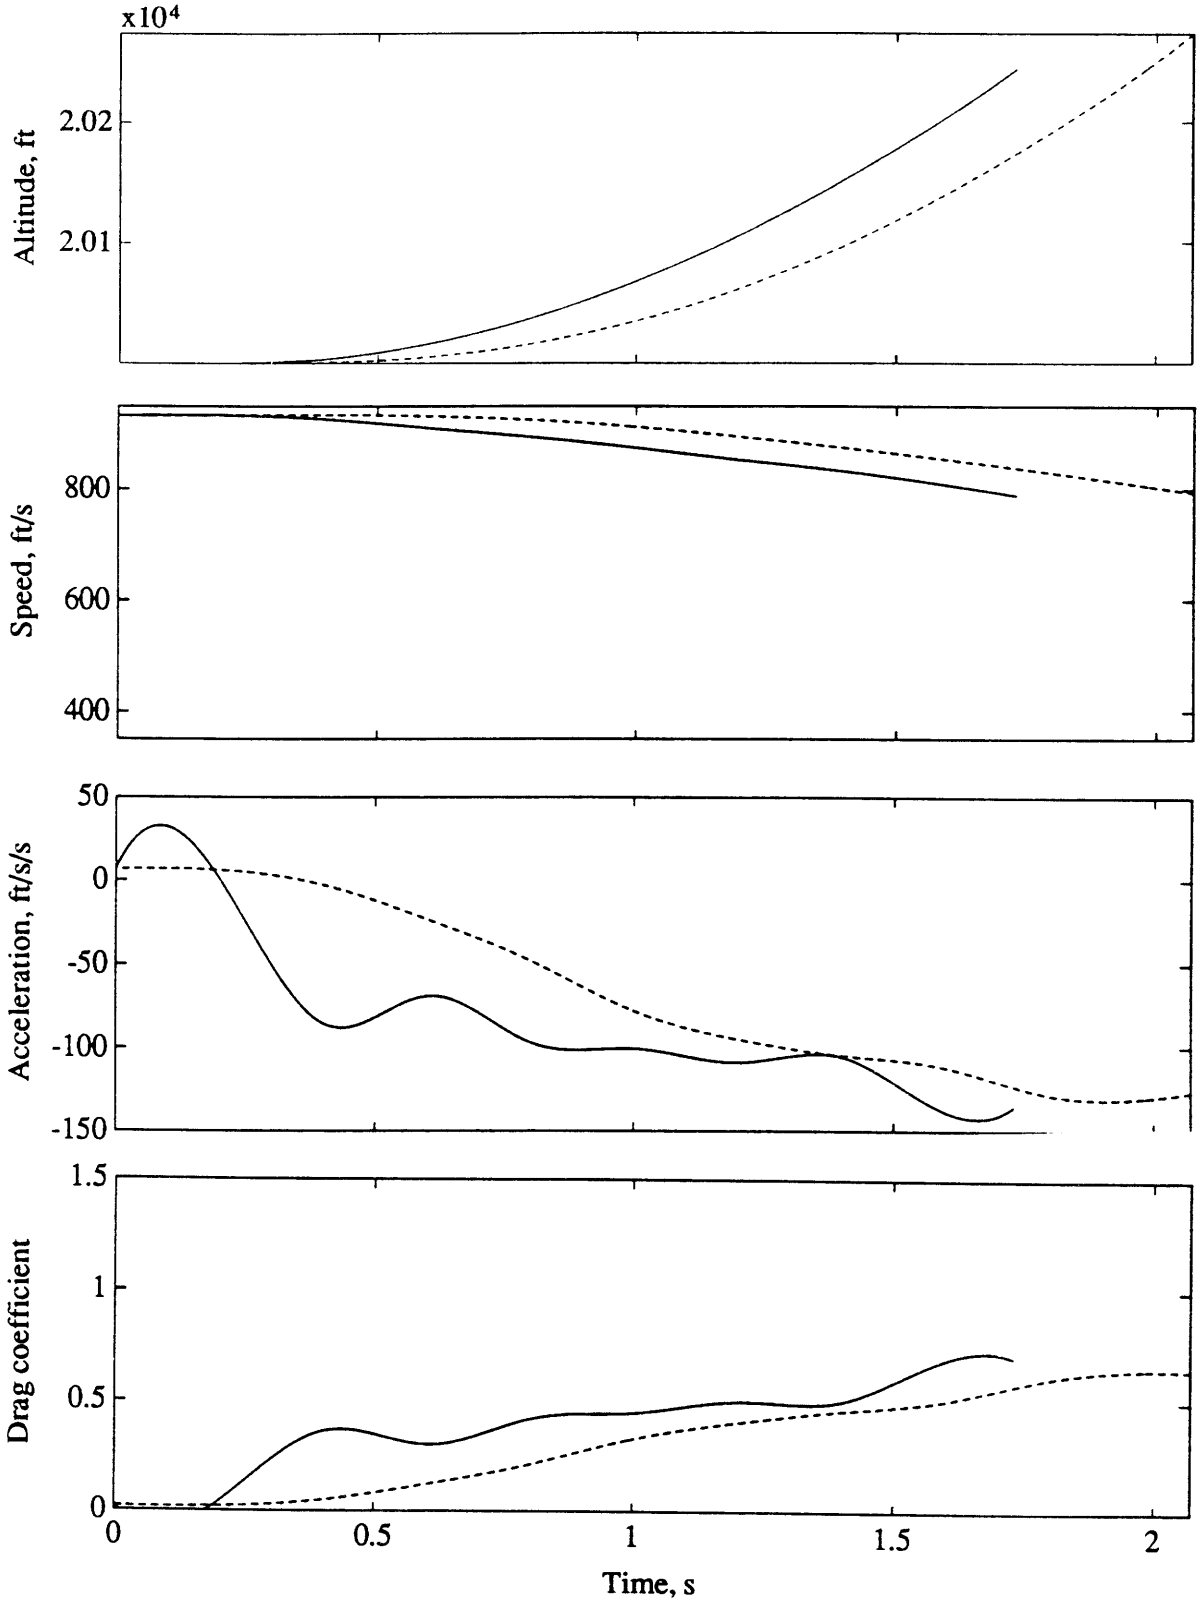

Figure 4.25 (cont'd)

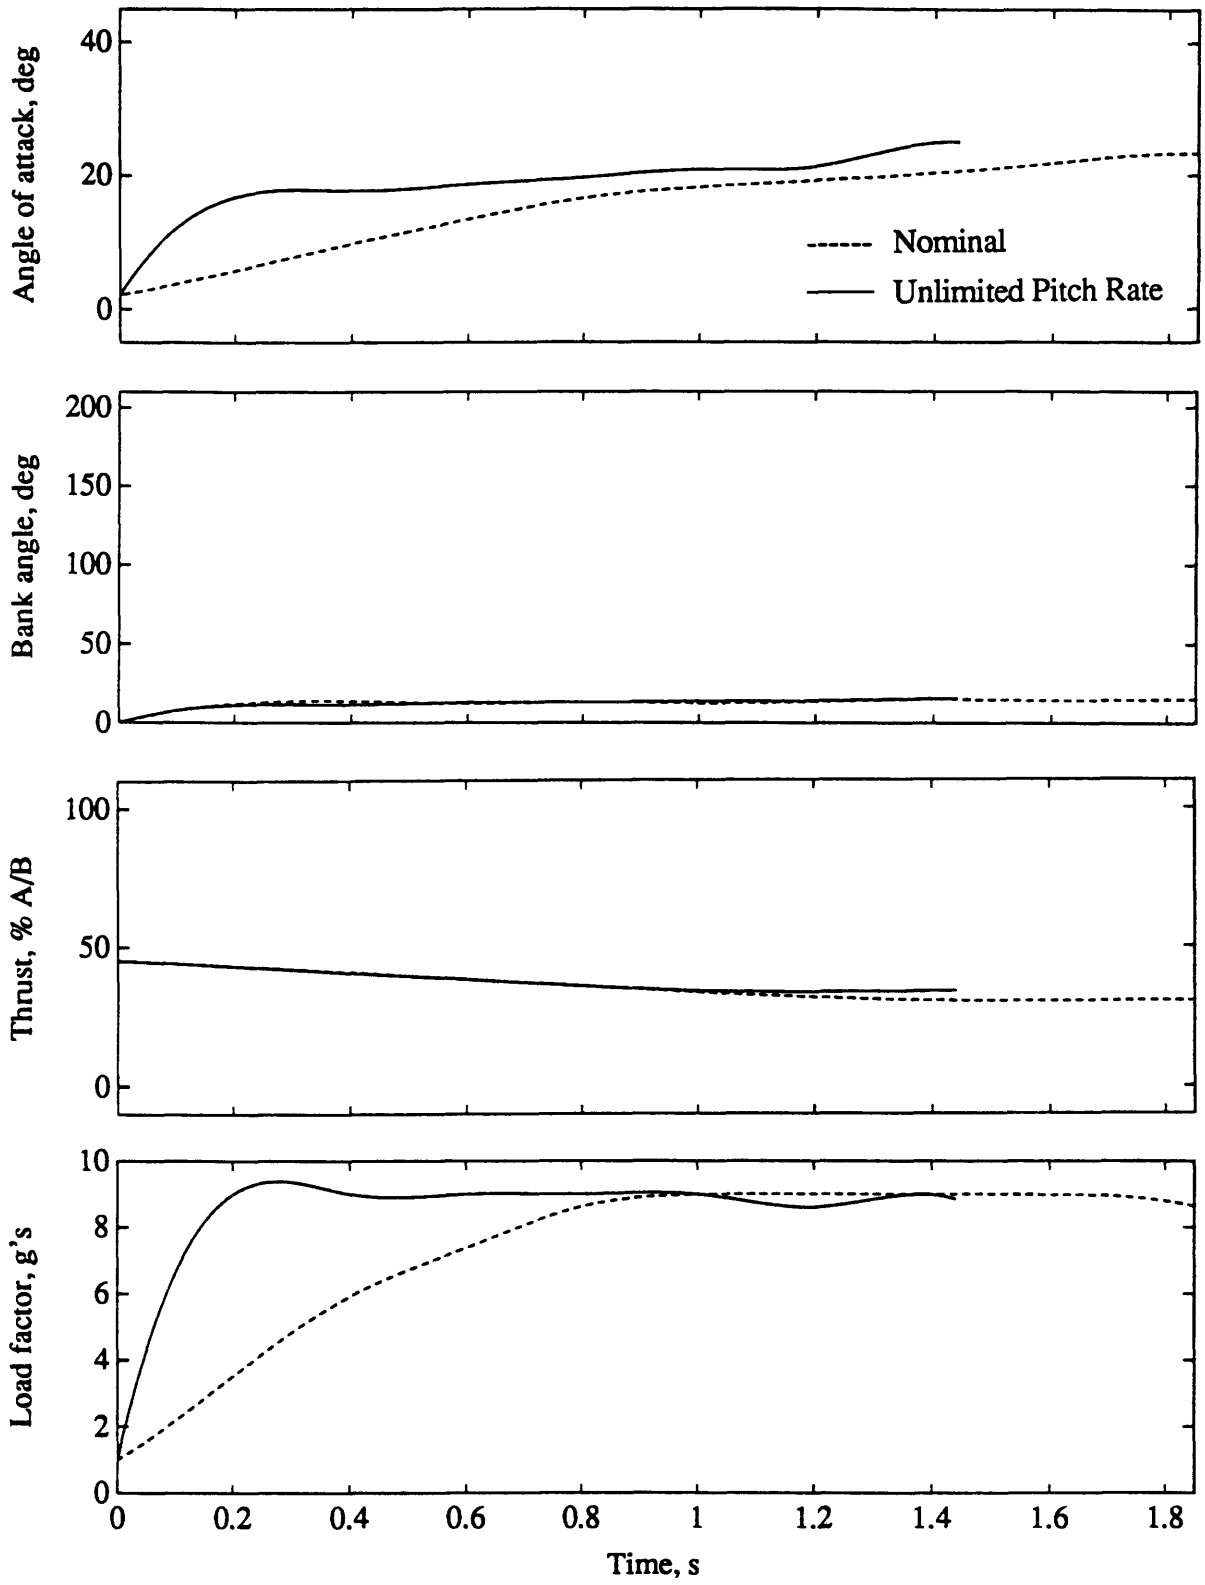

Figure 4.26. Optimal trajectories in 19 segments for nominal model vs. unlimited pitch rate model to (15,45). Compare with Figure 4.25, the same target in 9 segments. The unlimited α and load factor histories look more like steps, and both nominal and unlimited model times decreased.

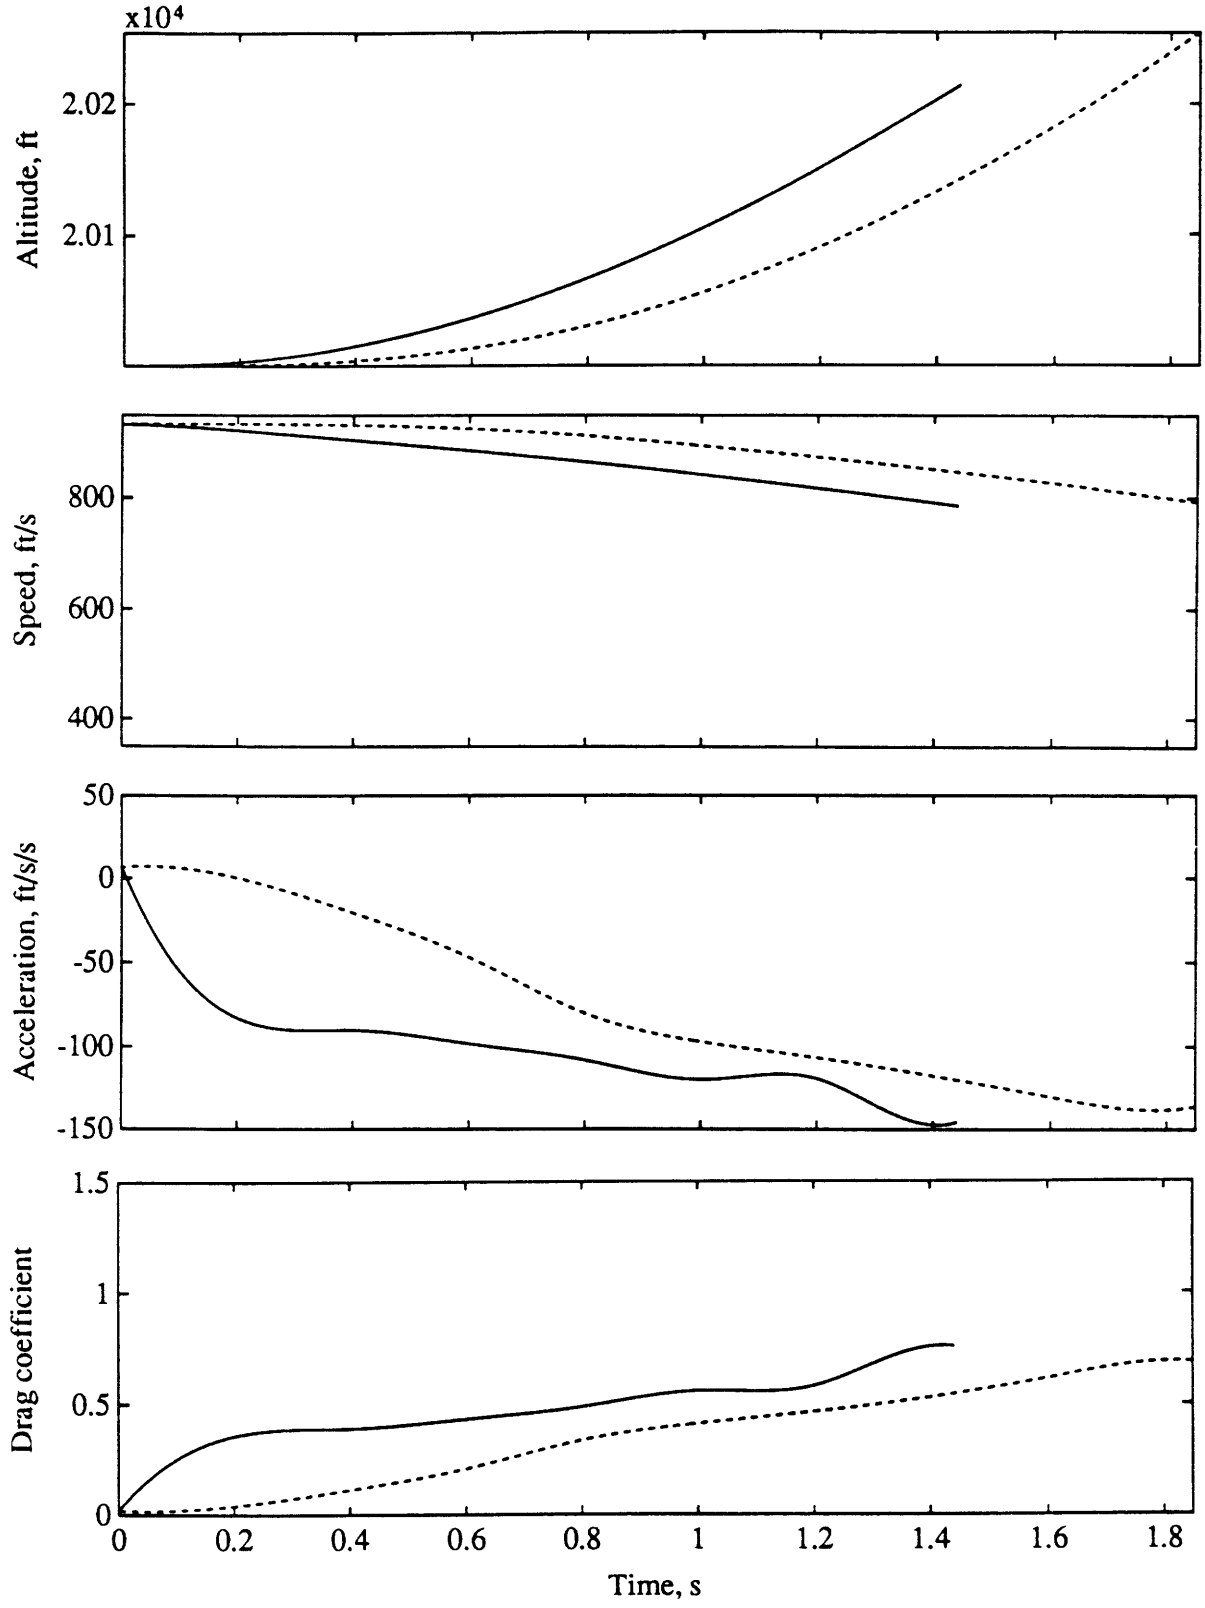

Figure 4.26 (cont'd)

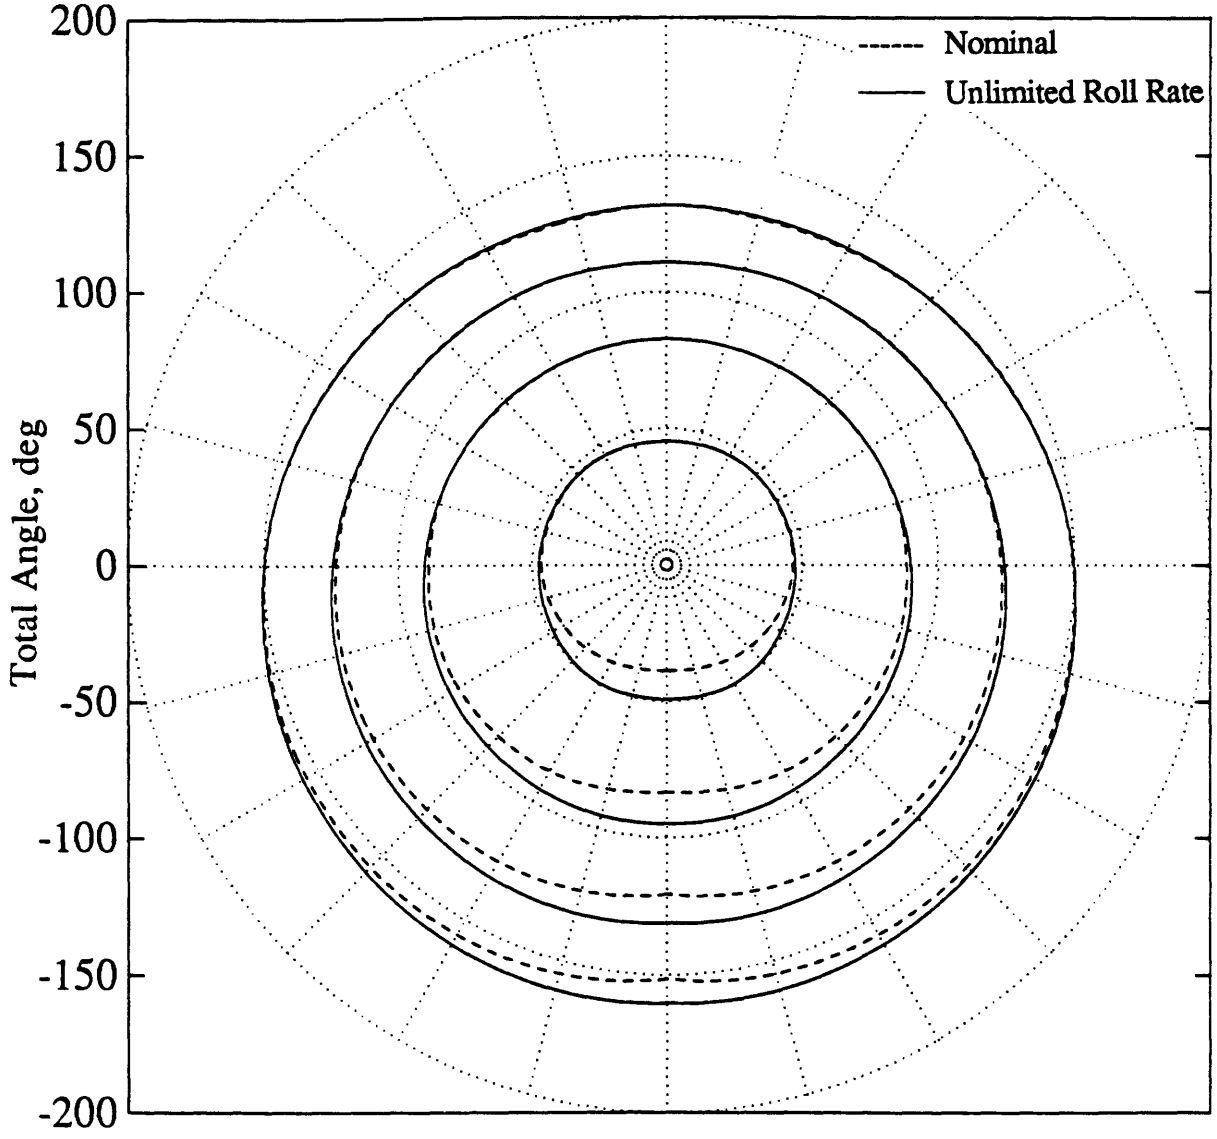

Figure 4.27. Contour plot of optimal trajectories of nominal model vs. unlimited roll rate model. The only time advantage comes in lower hemisphere pointing where the aircraft must roll more than 90 deg.

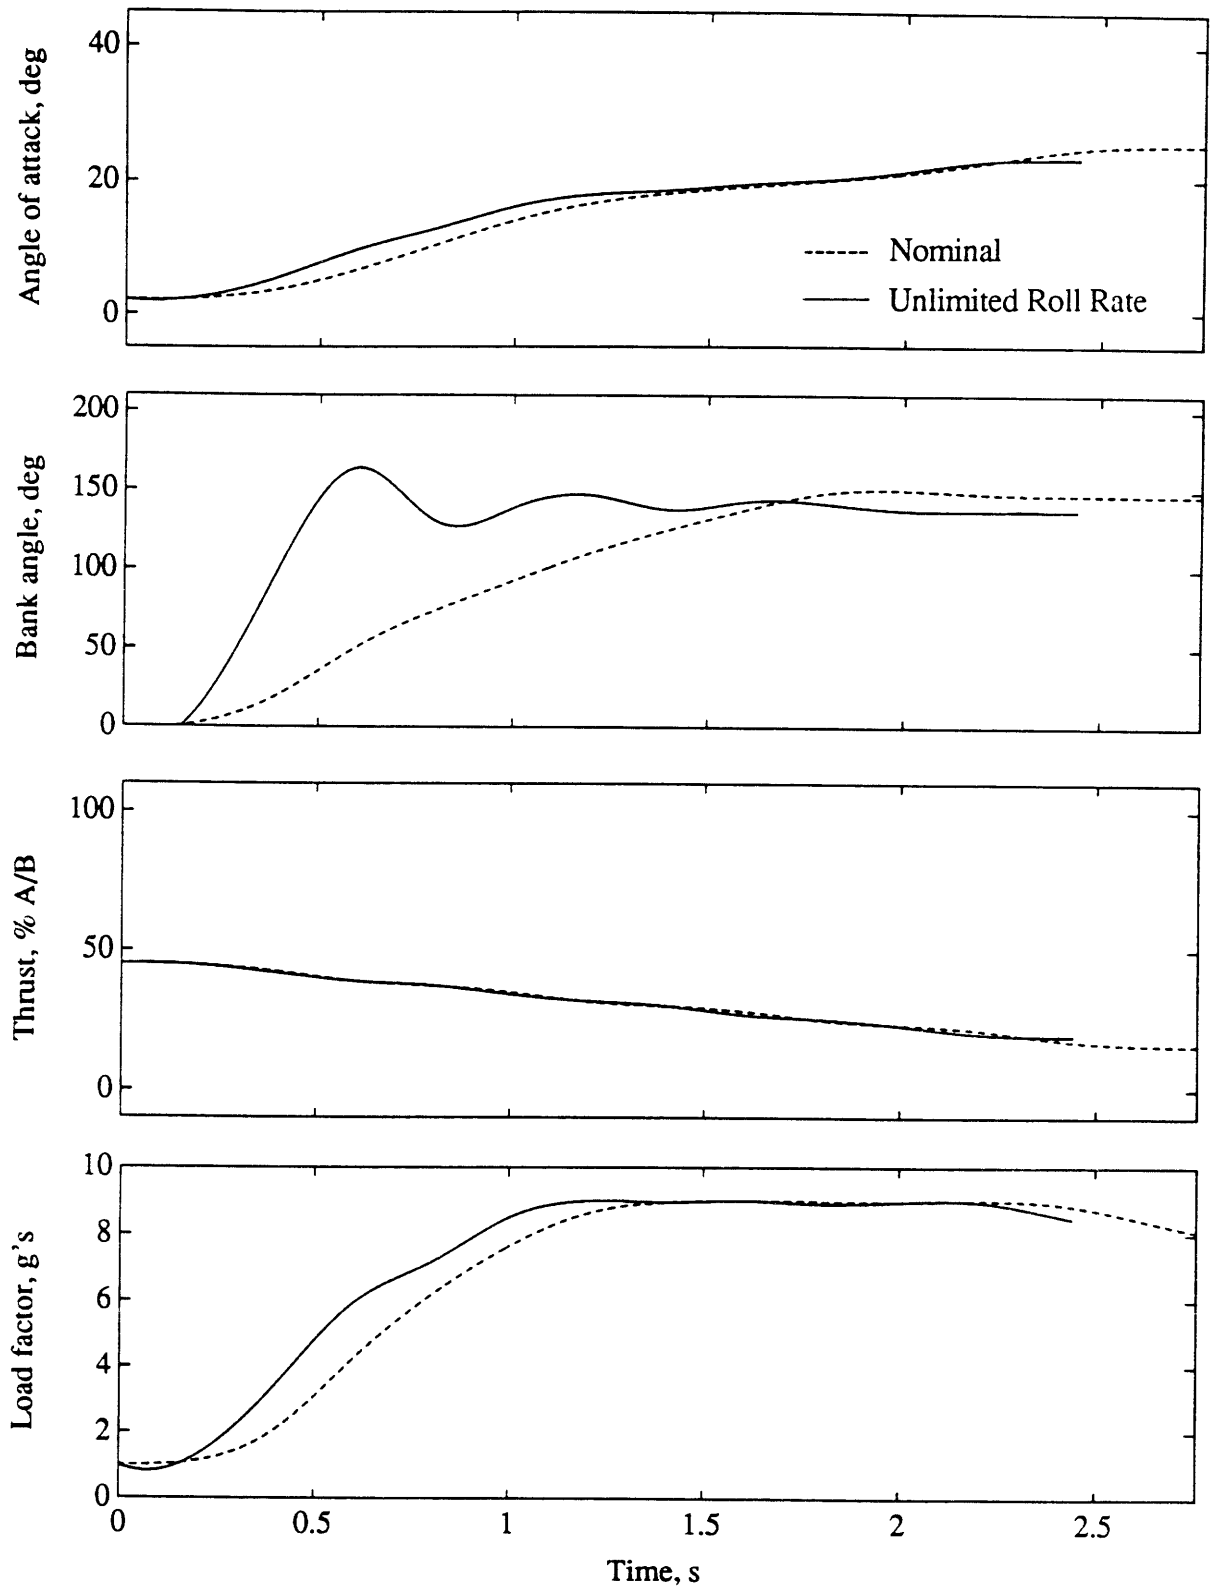

Figure 4.28. Optimal trajectories for nominal model vs. unlimited roll rate model to (45,-45). The required bank angle is reached much sooner for the unlimited model, so the aircraft may begin rolling sooner.

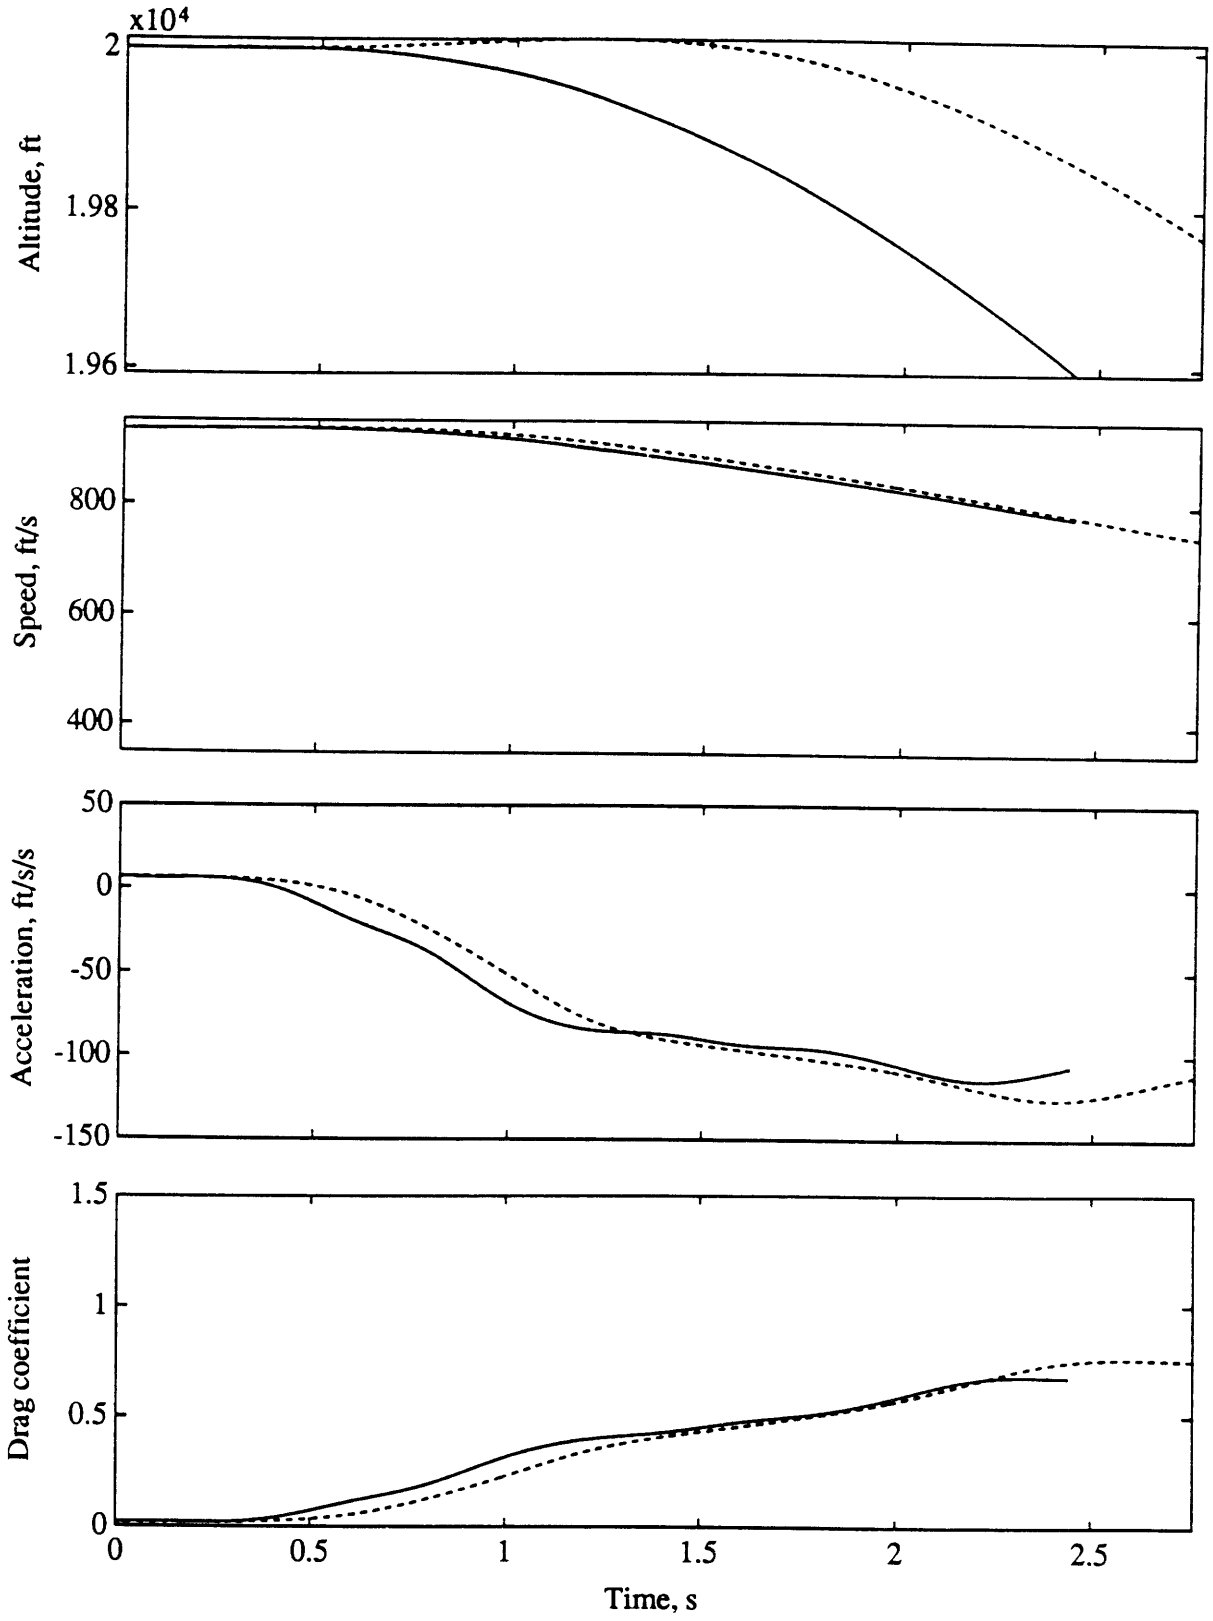

Figure 4.28 (cont'd)

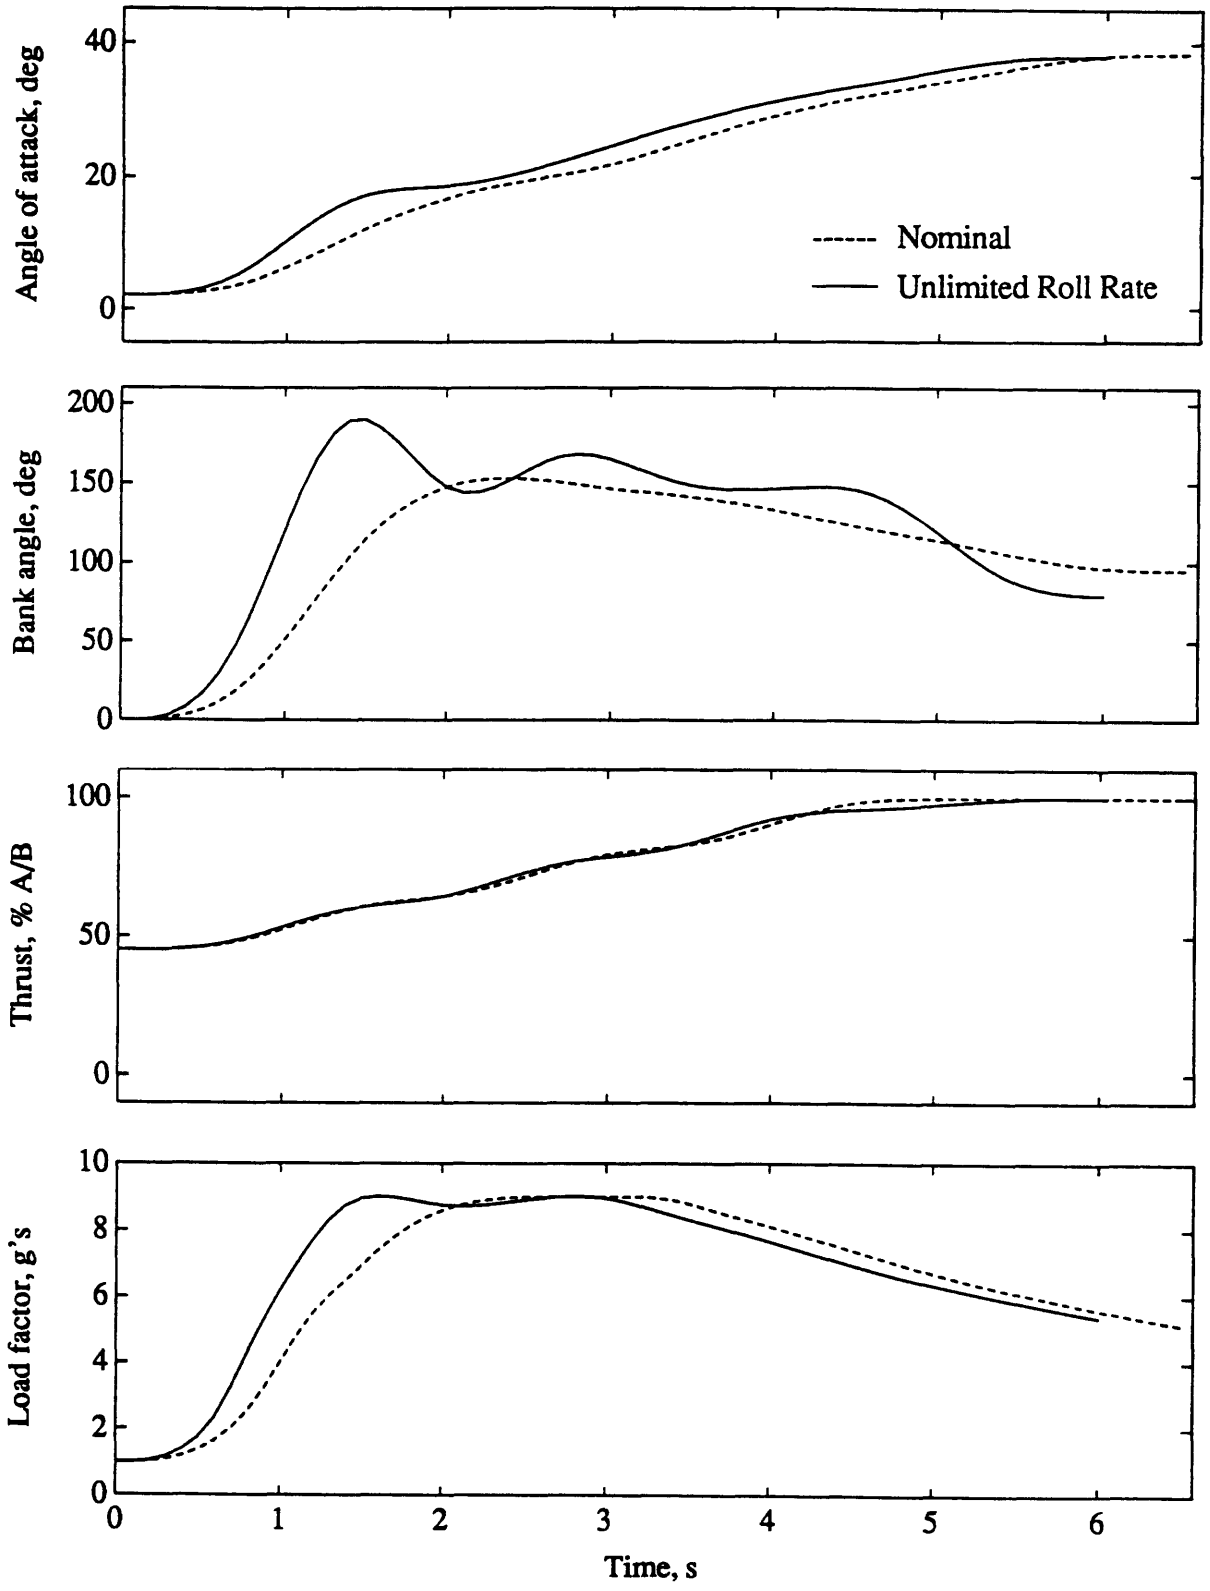

Figure 4.29. Optimal trajectories for nominal model vs. unlimited roll rate model to (165,-45). Note the oscillations in bank angle for the unlimited case, typical in lower hemisphere pointing.

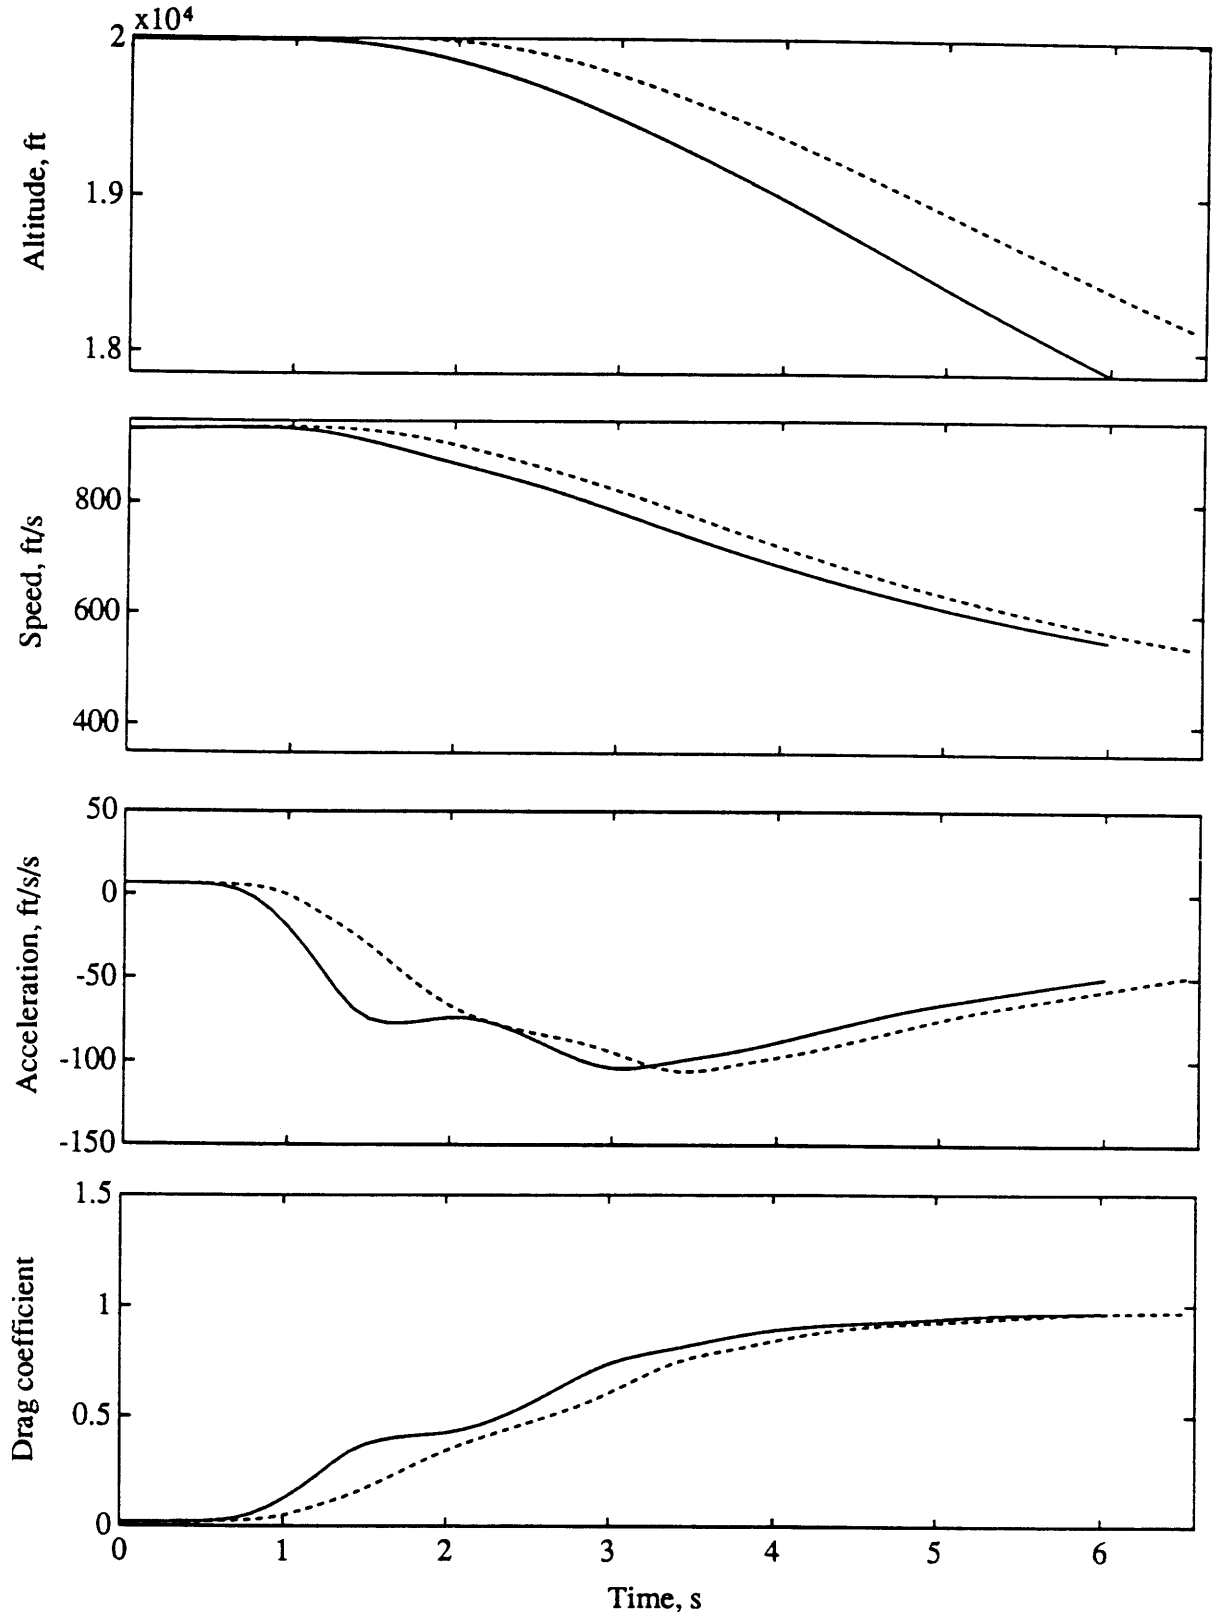

Figure 4.29 (cont'd)

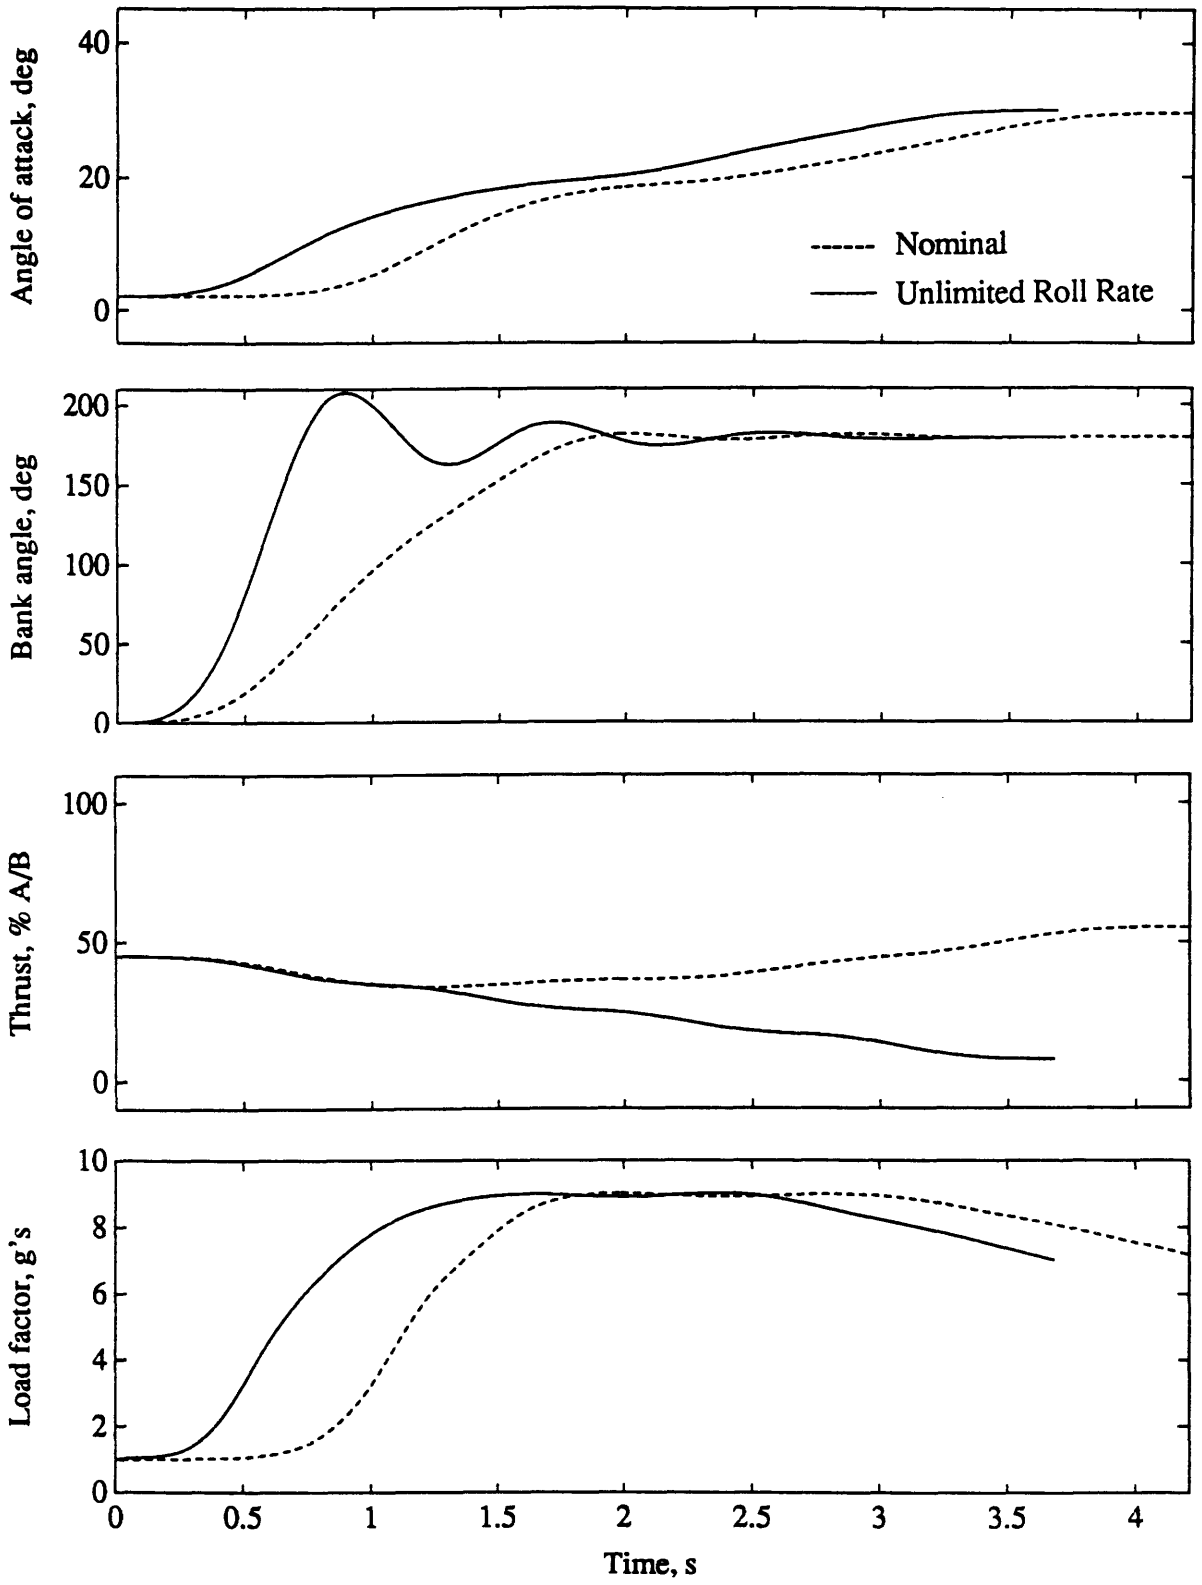

Figure 4.30. Optimal trajectories for nominal model vs. unlimited roll rate model to (15,-88). Thrust is different for the two models so the unlimited model can increase α faster without violating the load factor limit.

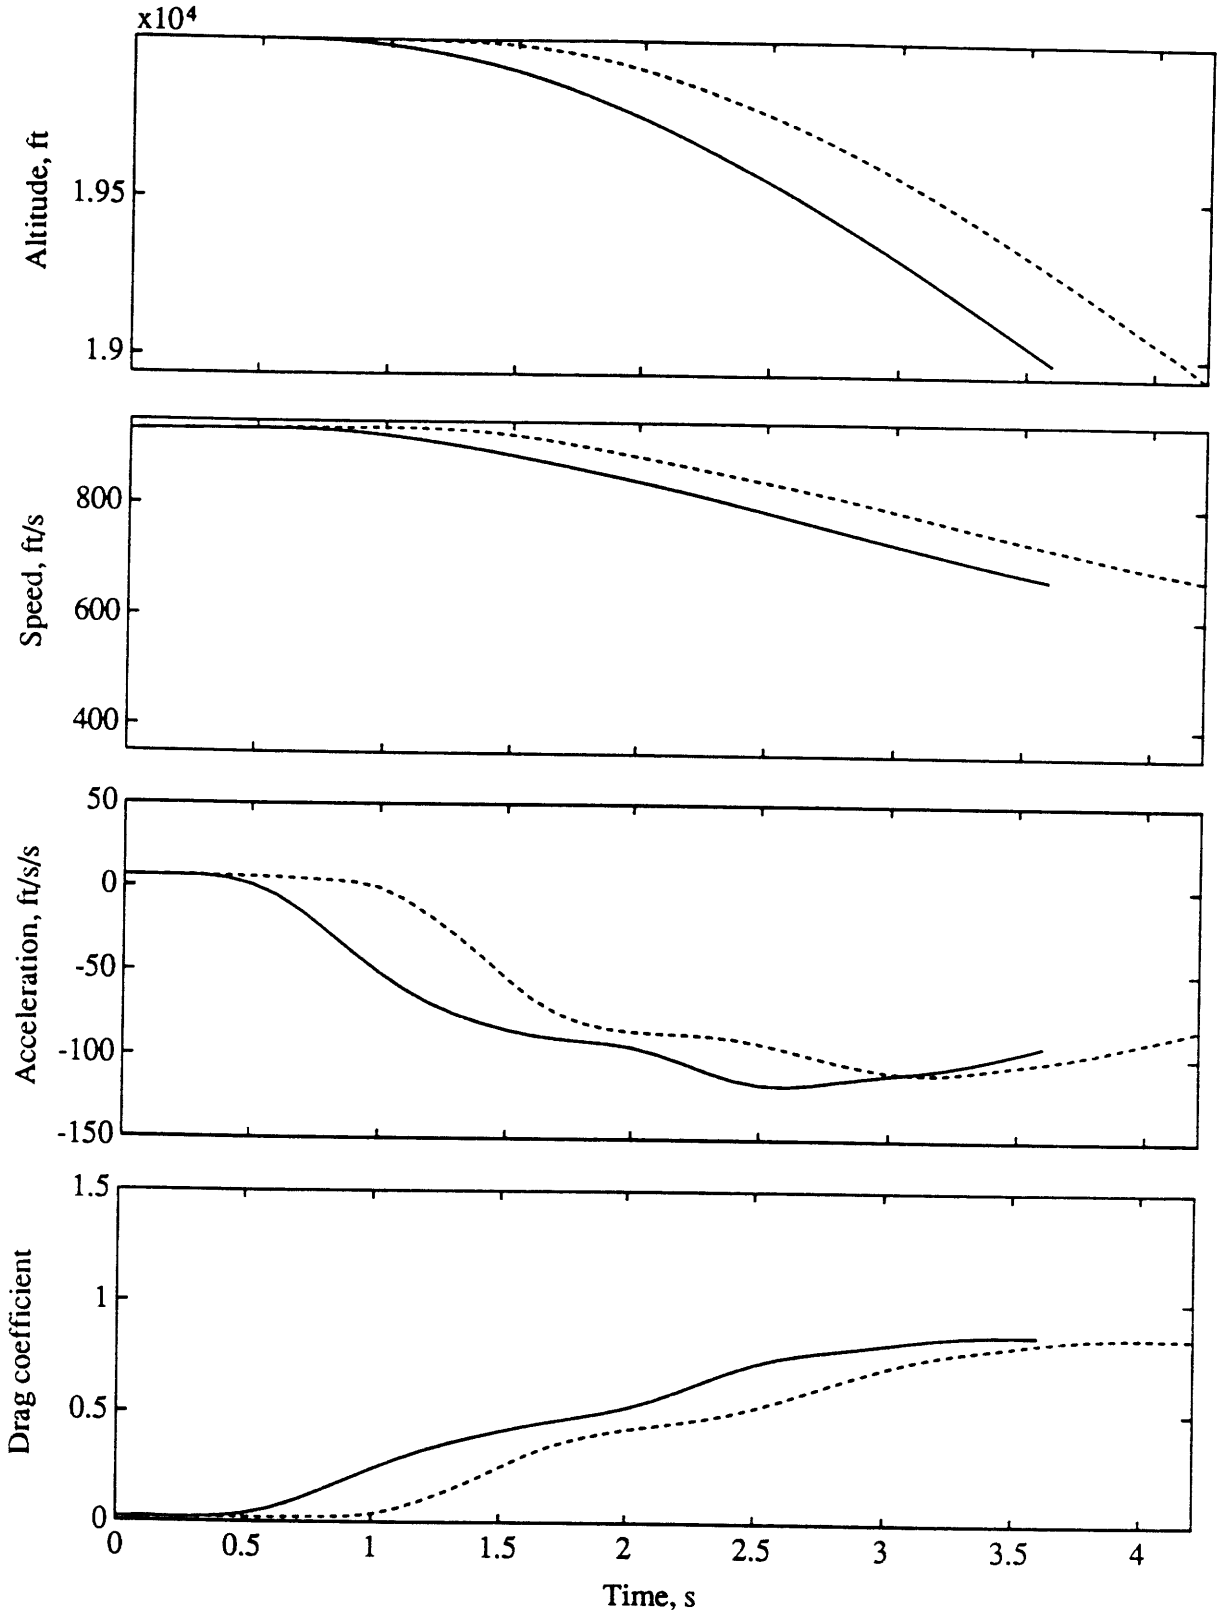

Figure 4.30 (cont'd)

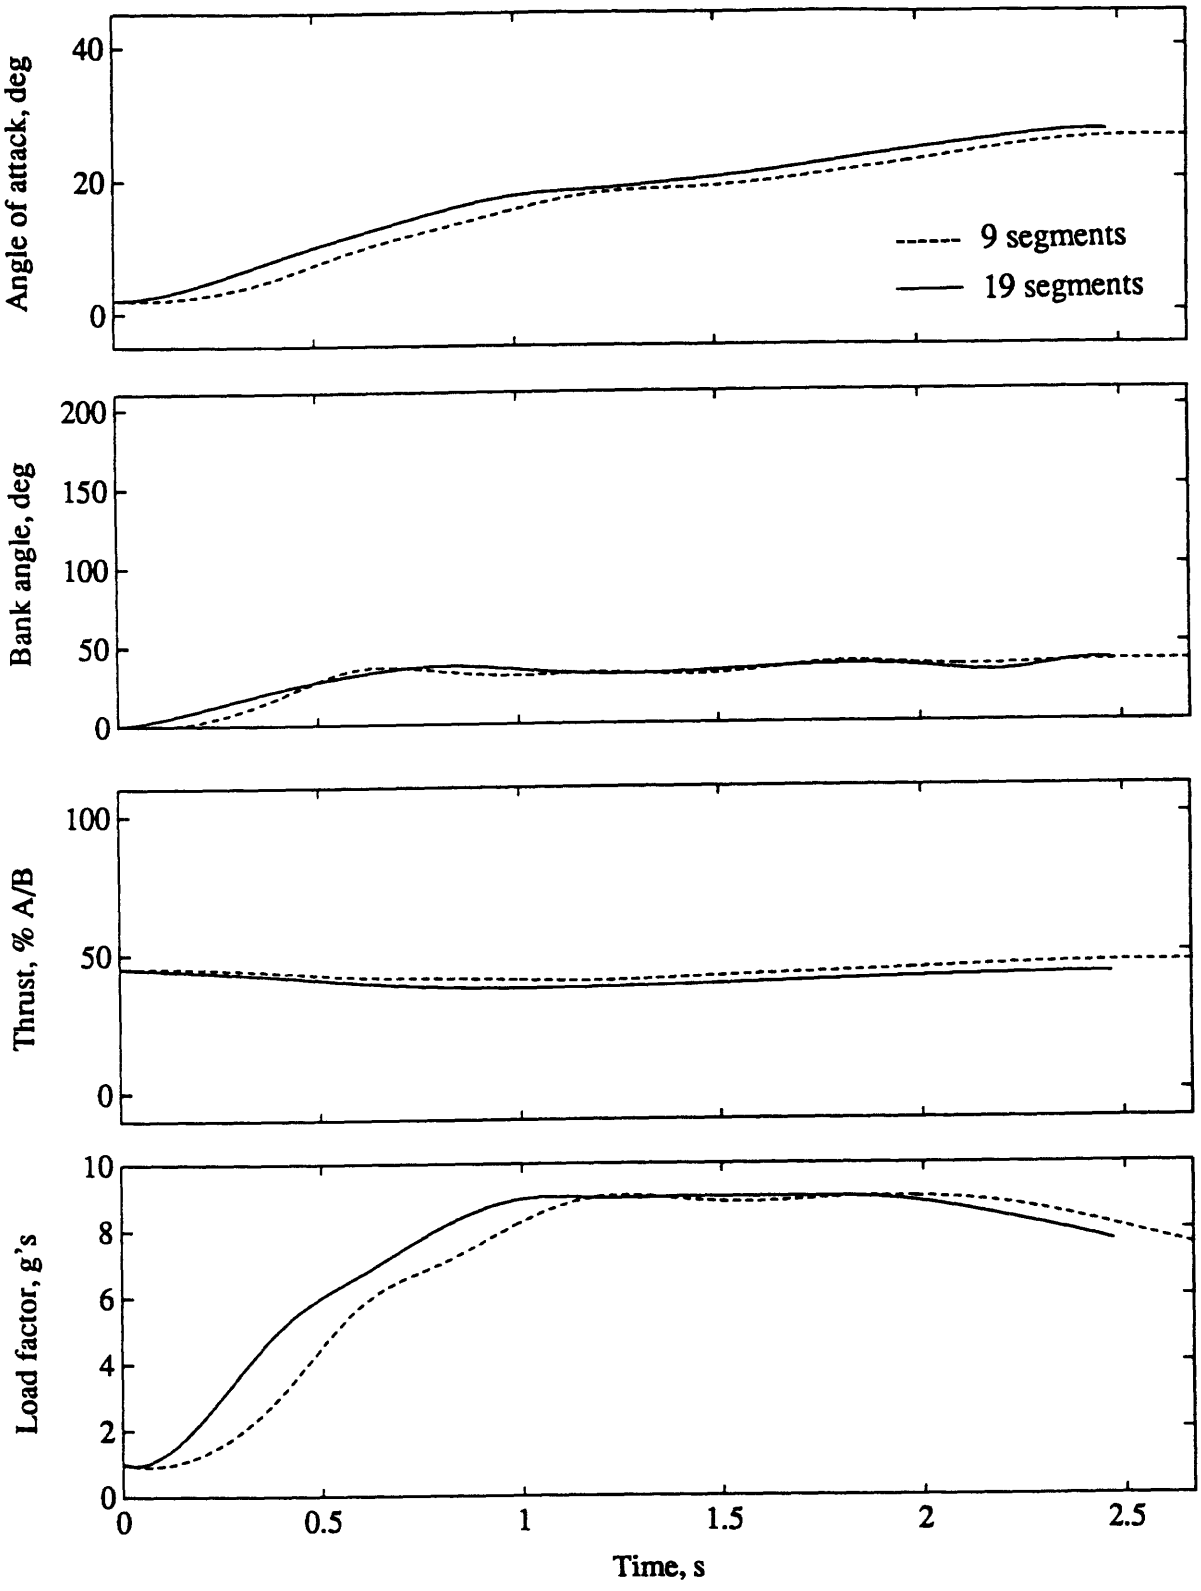

Figure 4.31. Optimal trajectories for unlimited roll rate model to (45,45) in 9 and 19 segments. The bank angle rises faster in the 19 segment case, dropping .2 s from the time in the process.

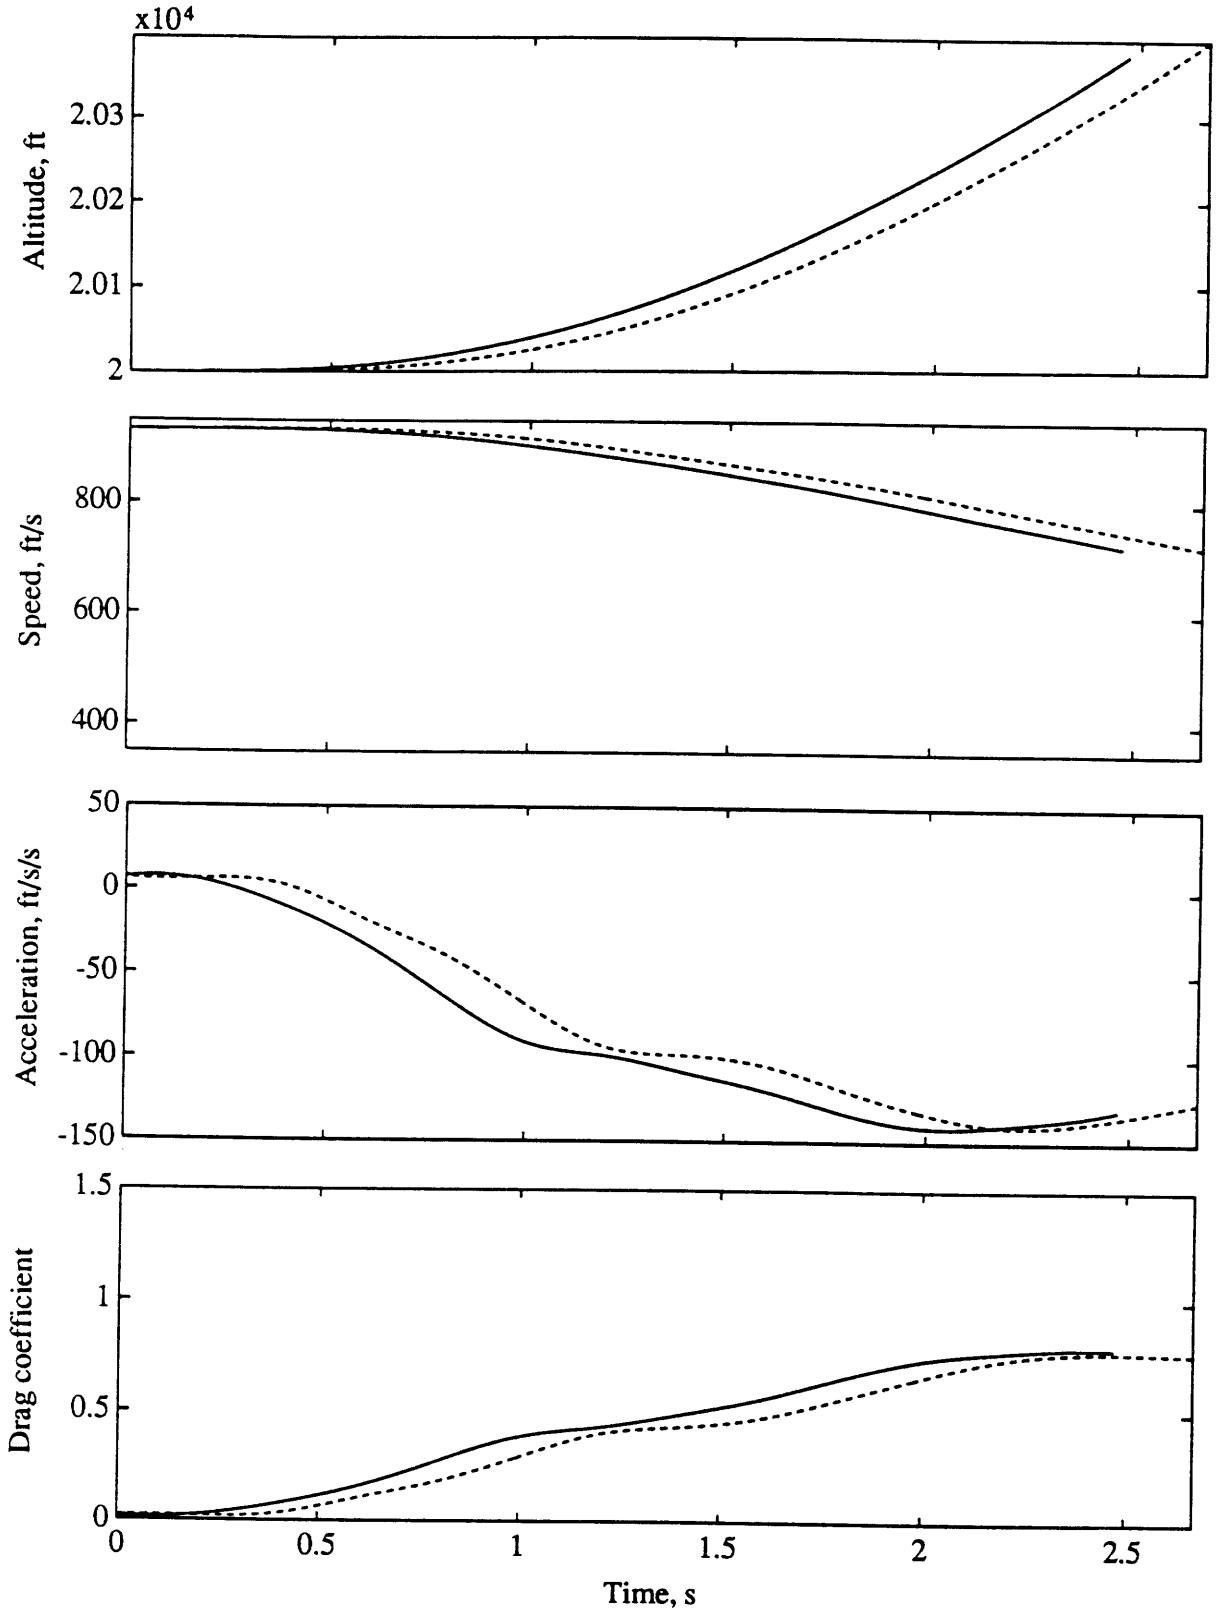

Figure 4.31 (cont'd)

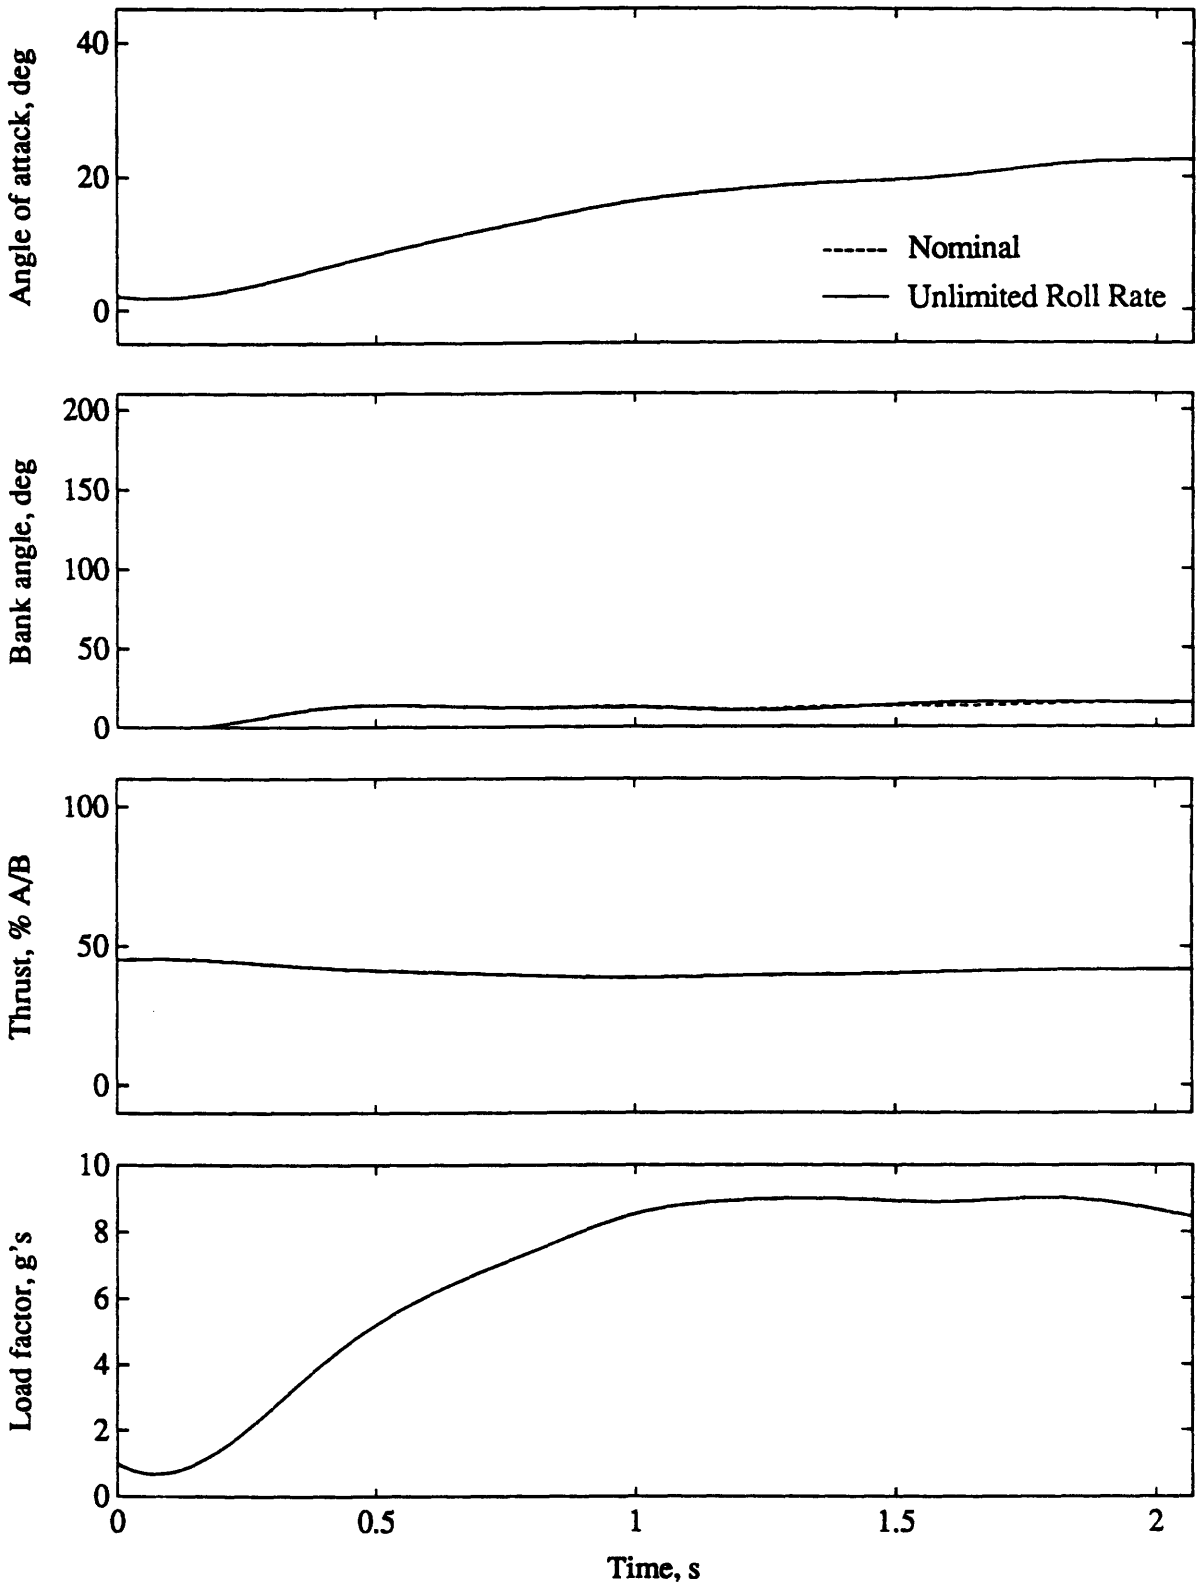

Figure 4.32. Optimal trajectories for nominal model vs. unlimited roll rate model to (15,45). Very little bank angle is required, so there is no advantage in removing the roll rate limits. The histories are virtually identical.

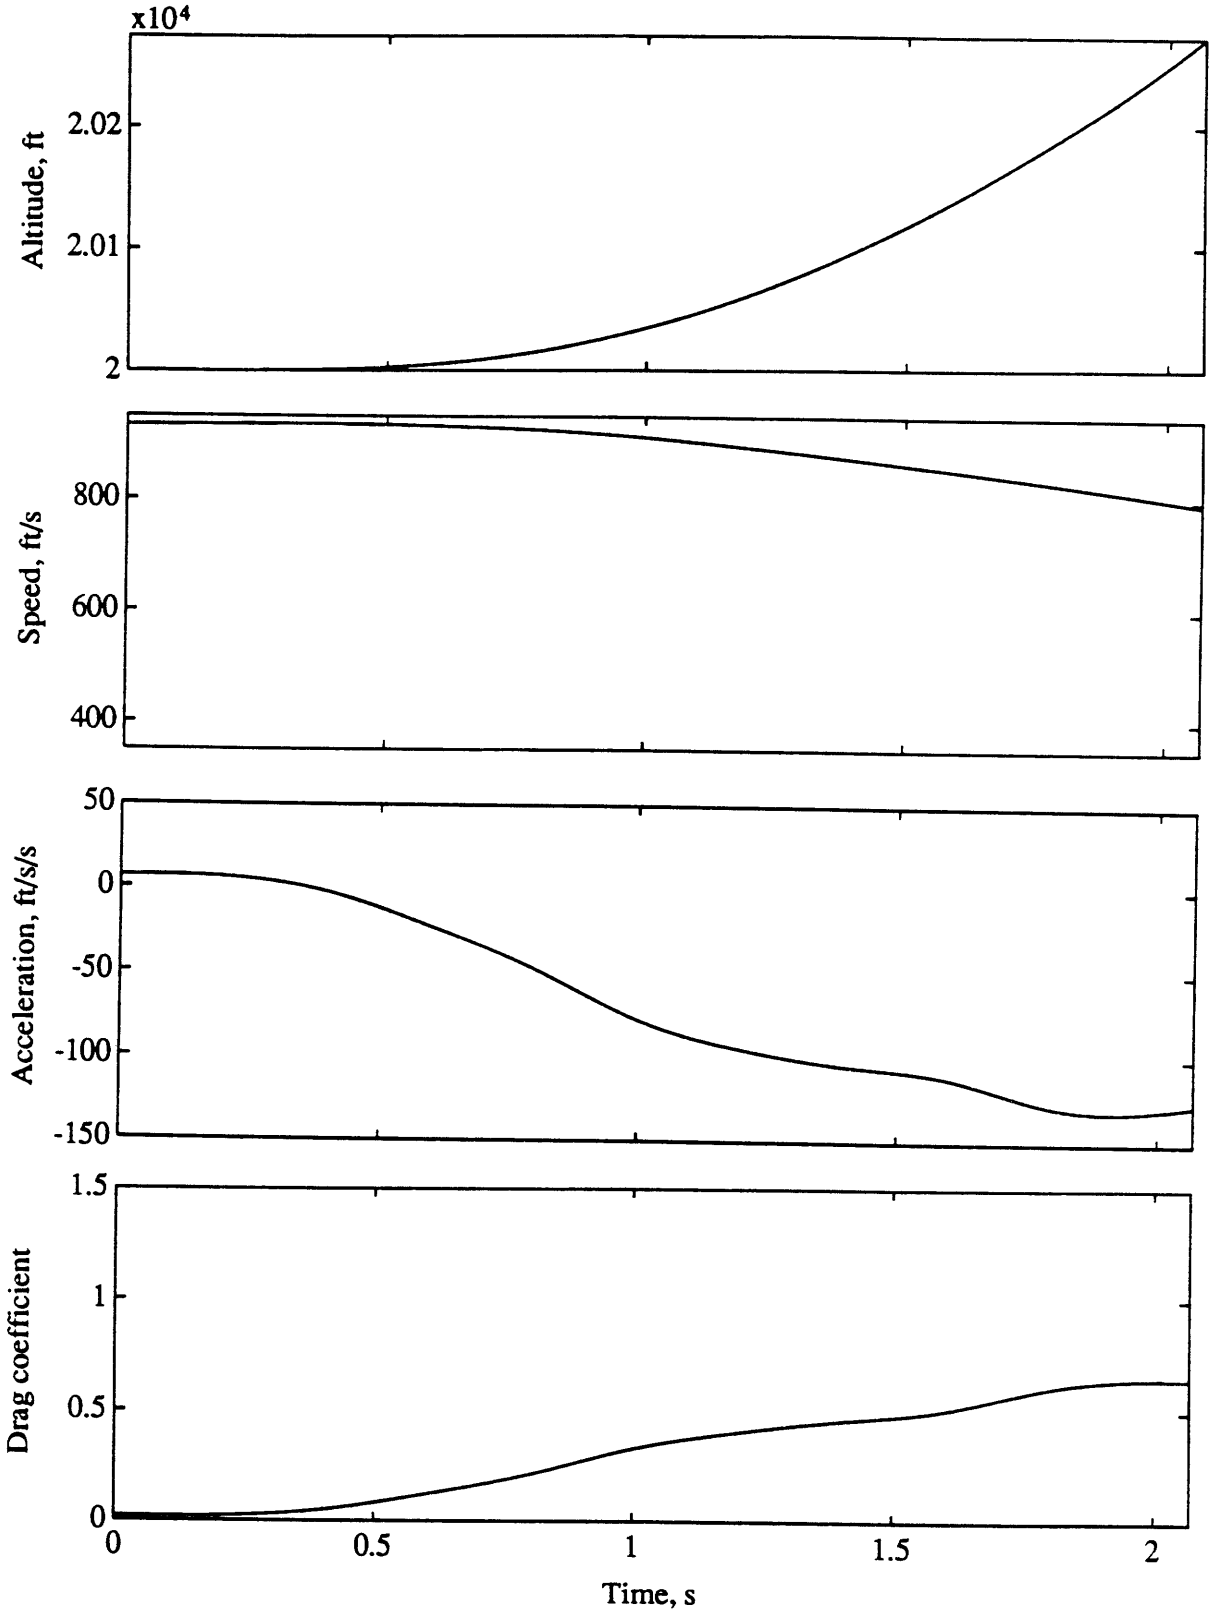

Figure 4.32 (cont'd)

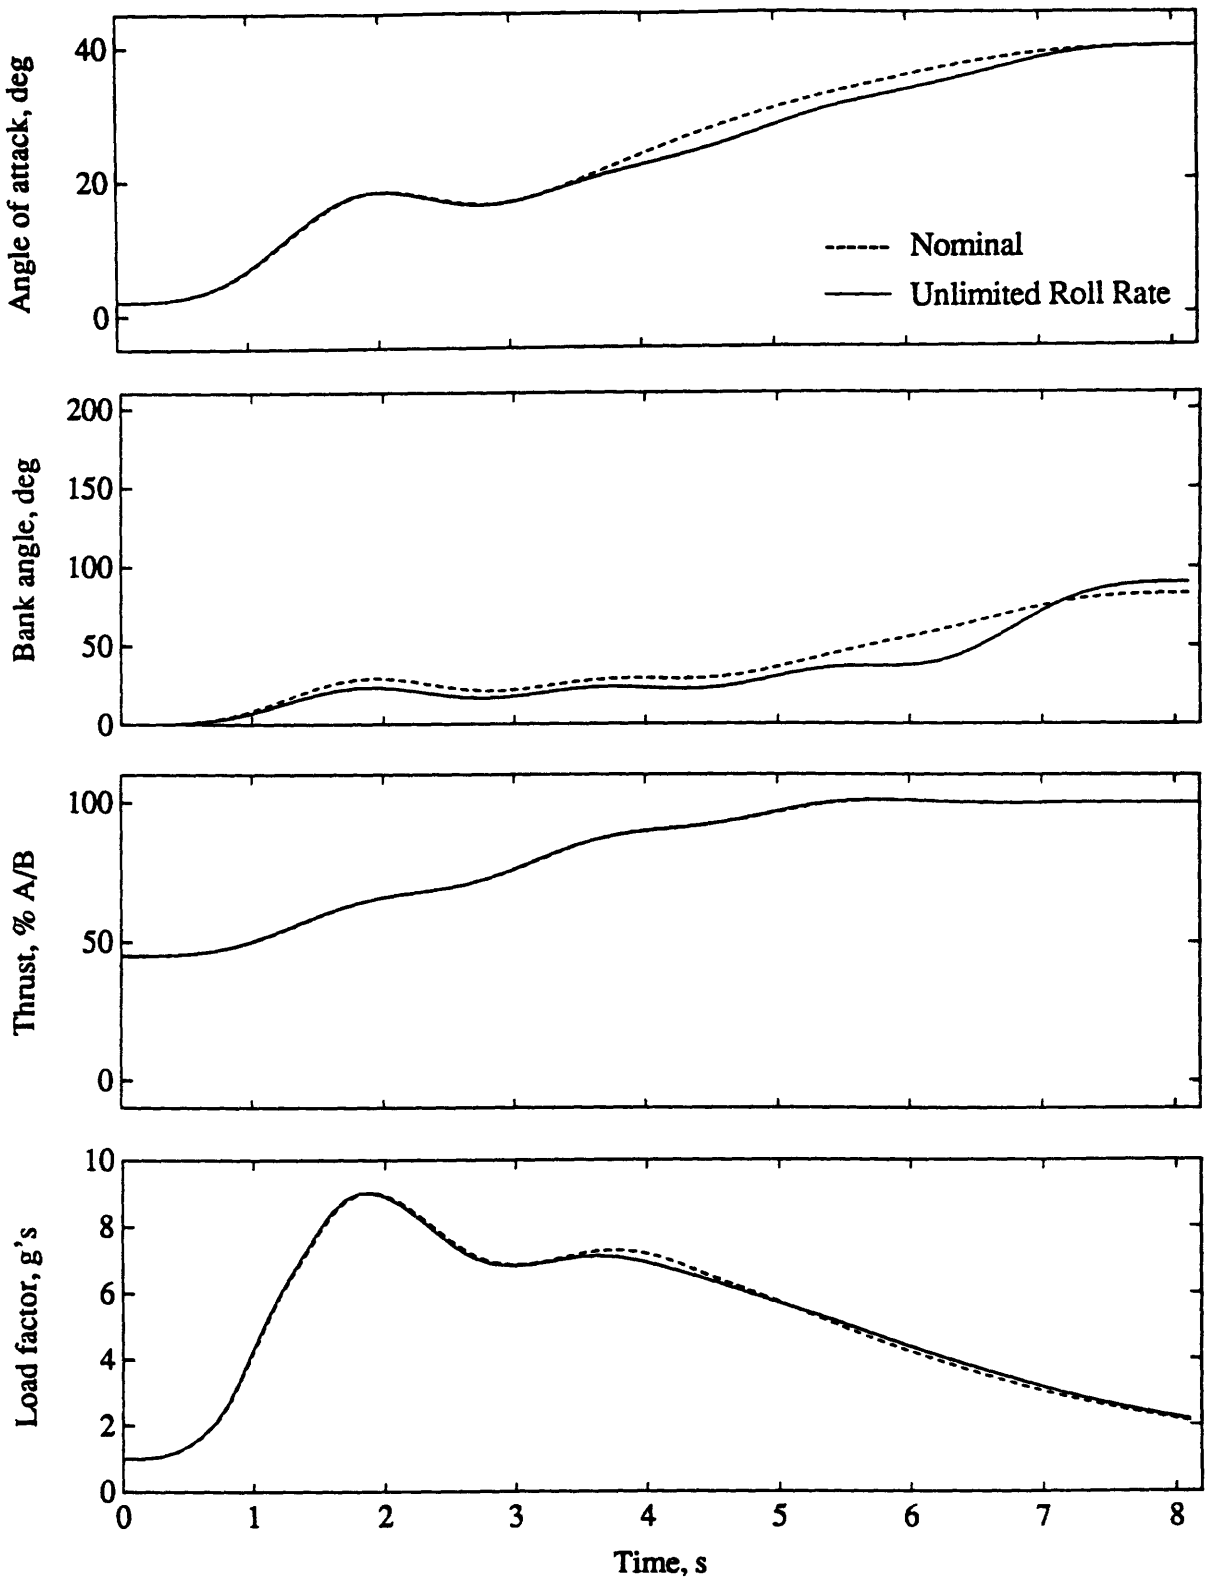

Figure 4.33. Optimal trajectories for nominal model vs. unlimited roll rate model to (165,45). The unlimited bank angle is actually less than the nominal.

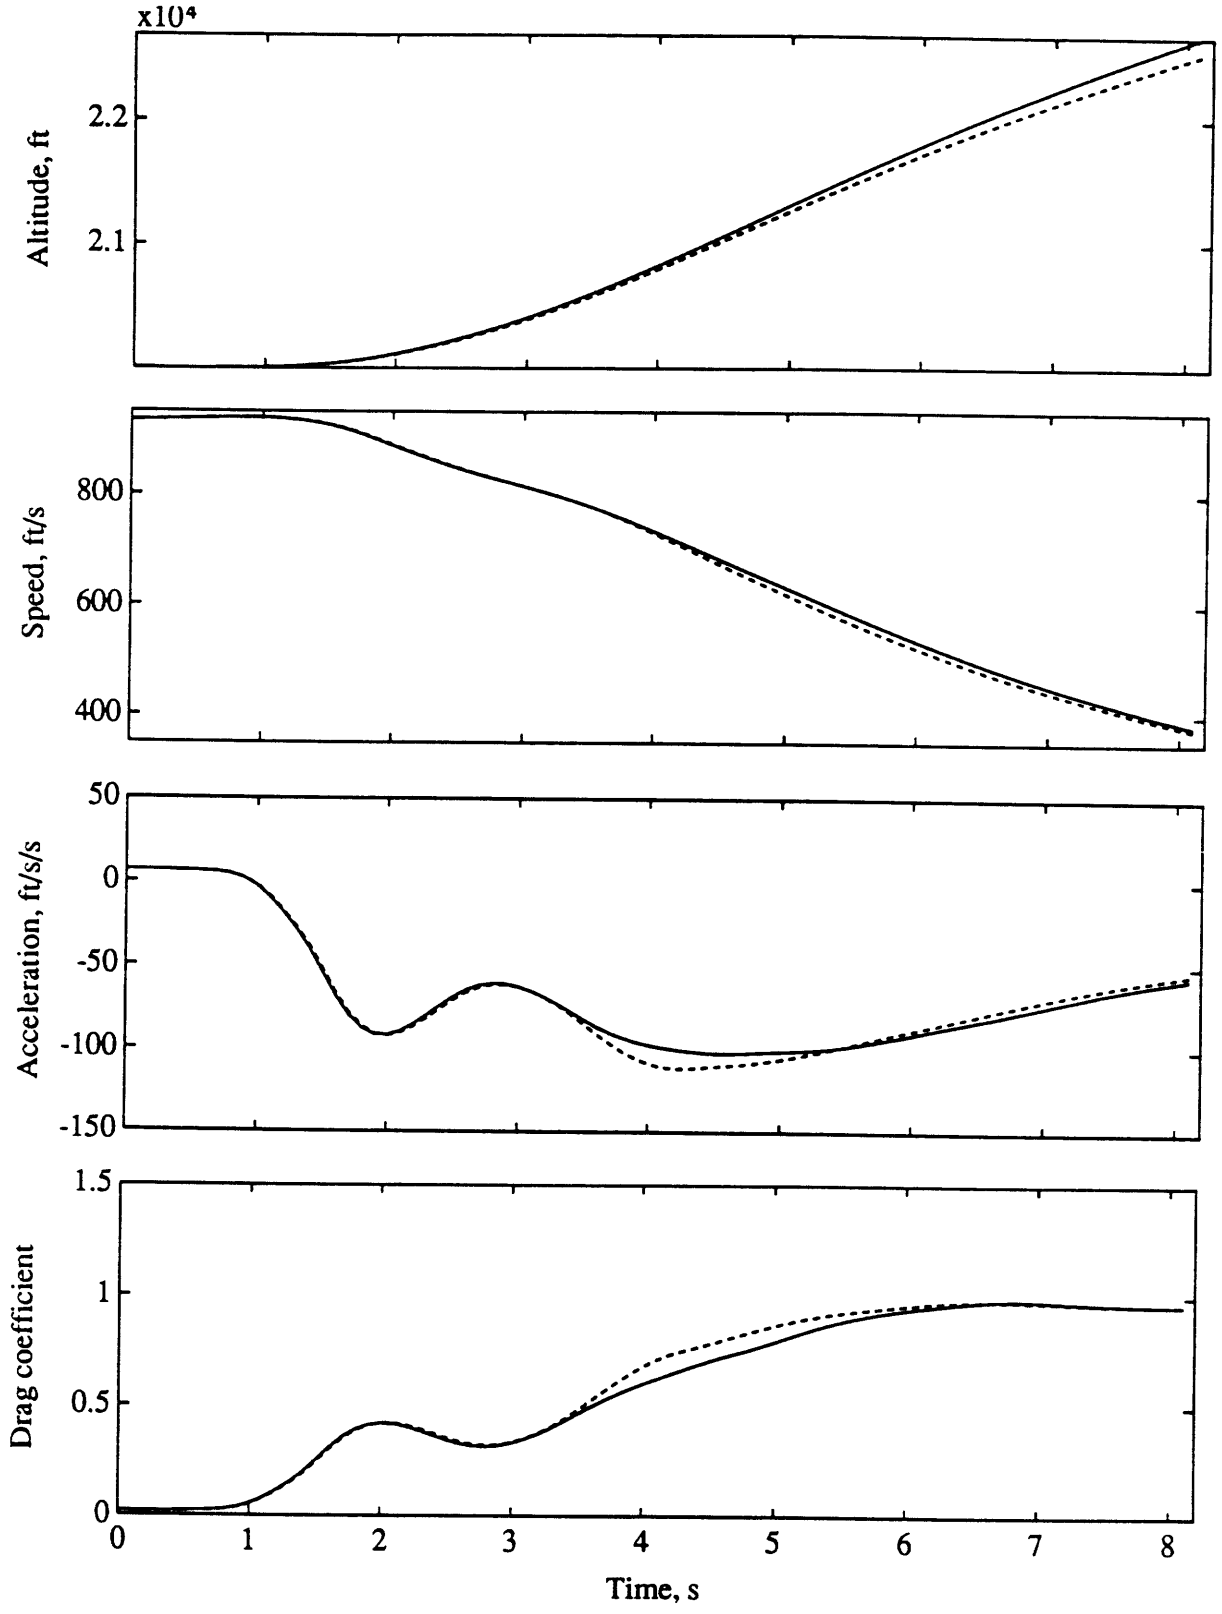

Figure 4.33 (cont'd)

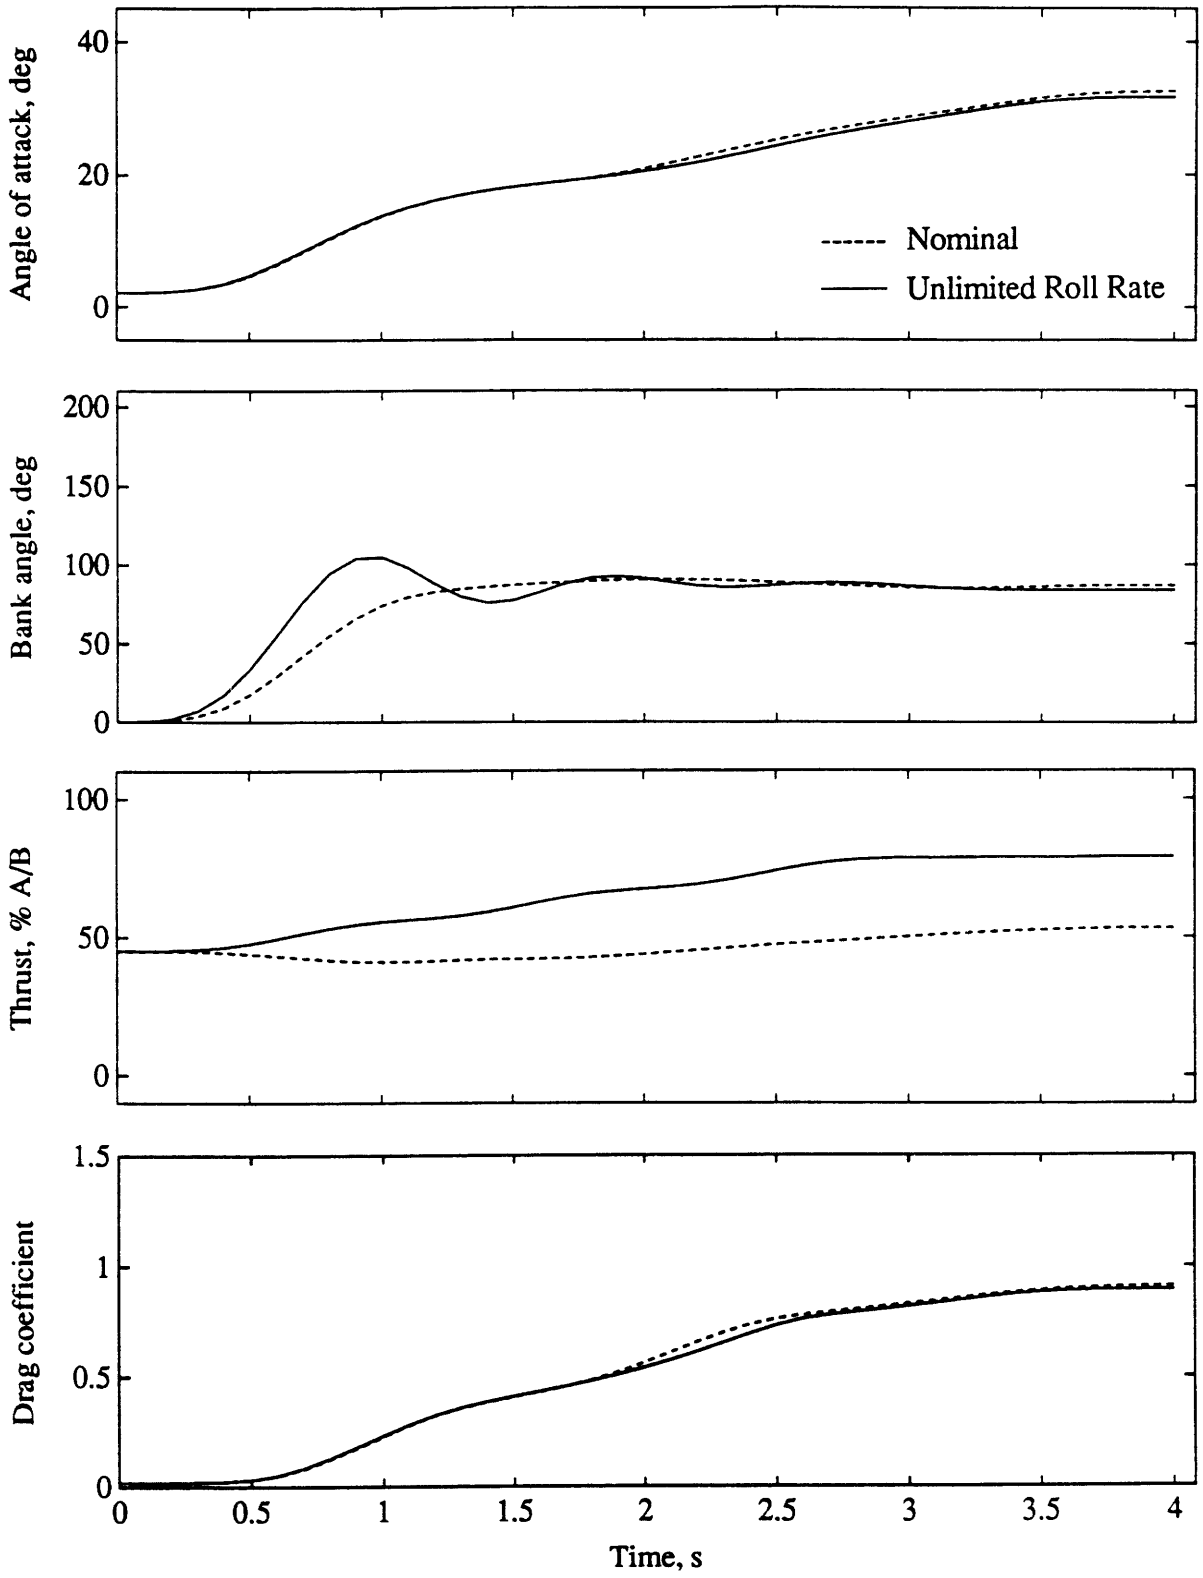

Figure 4.34. Optimal trajectories for nominal model vs. unlimited roll rate model to (90,0). The unlimited thrust is higher to keep speed and load factor up. Too high of an angle of attack is not desired because the airplane is pulled up from its target elevation.

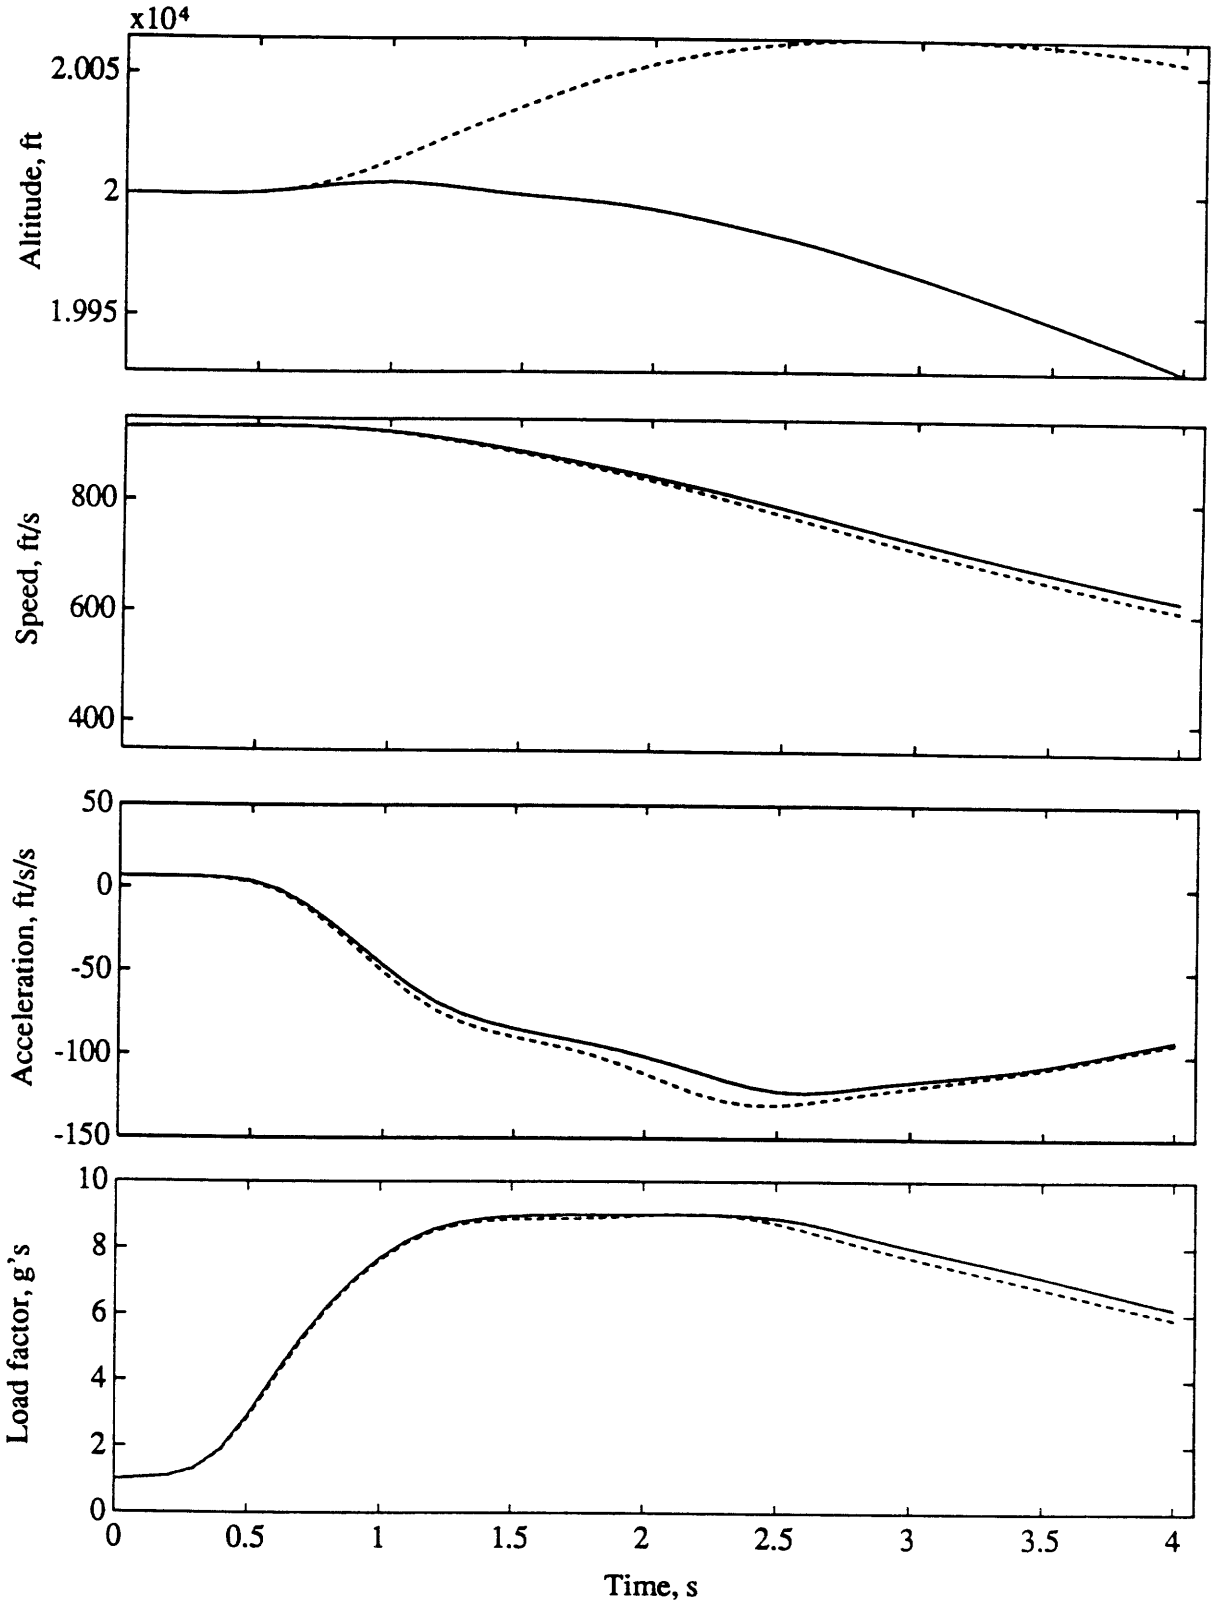

Figure 4.34 (cont'd)

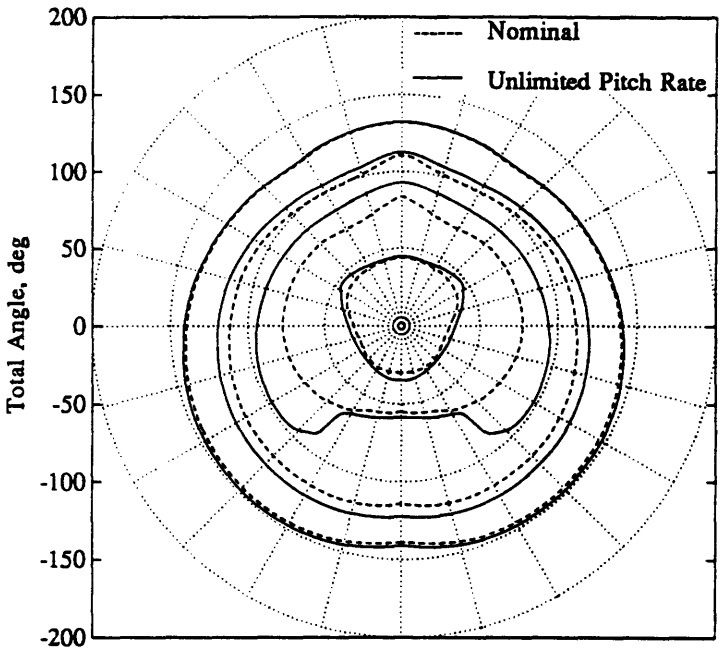

(a)

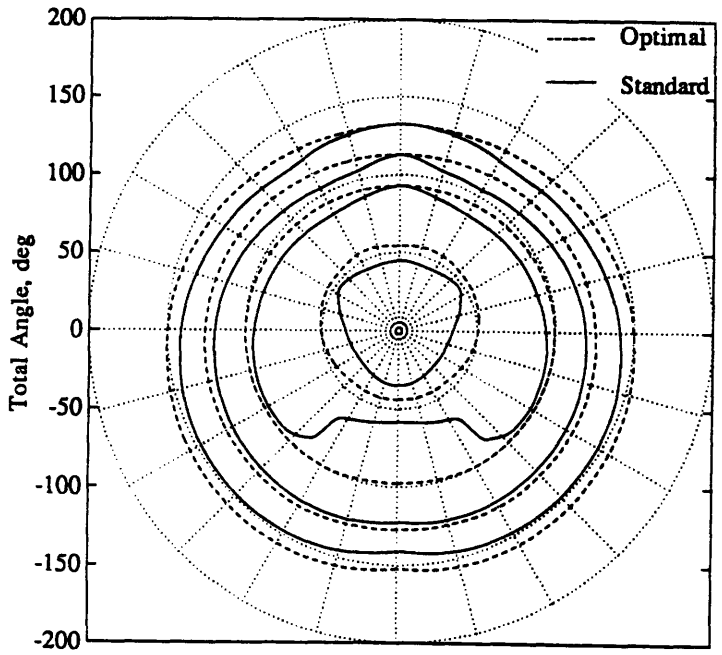

(b)

Figure 4.35 Agility comparison using standard maneuvers. (a) shows the contour plot of standard maneuvers for the nominal and unlimited pitch rate models. (b) compares standard and optimal trajectories for the unlimited pitch rate model.

5 Conclusions and Recommendations

5.1 Conclusions

In this thesis, agility was studied using optimal control. The goals of the research were to:

1. Demonstrate that faster pointing can be achieved by optimizing aircraft trajectories.
2. Discover the effects of performance limits on the optimal trajectories.
3. Demonstrate that using standard maneuvers to compare aircraft agility may hide true performance capabilities.

These goals were met by simulating six sets of pointing maneuvers to various heading and elevation angle combinations for a high-performance fighter model. Two sets used a nominal model with three performance limitations in standard and time-optimal maneuvers. Three more sets were the optimal trajectories for the model with a different performance limitation removed in each set. The final set was for an unlimited pitch rate model using standard maneuvers.

There is a significant time advantage, up to 48%, for the optimal trajectories to targets requiring lift vector reorientation. The edge exists because rolling and pitching were performed simultaneously in the optimal case. In the standard maneuvers, the aircraft can roll only at 1 g. For targets not requiring lift vector reorientation, that is, those with an elevation of 88 deg, there is no time advantage for the optimal case because no rolling is required. Similarly, the aircraft performs a split-S to the targets with a heading of 179 deg and an elevation of 60 deg or 45 deg instead of a bank-and-pull to eliminate the roll disadvantage. In these two cases, the split-S is a global minimum while the bank-and-pull is a local minimum.

Some of the performance limits are less detrimental than others in the pointing maneuver. Engine spool time had virtually no effect on pointing time: the largest time difference between the nominal and unlimited spool time runs was 0.2 s. The differences are magnified slightly for longer trajectories, indicating that energy management becomes more important for sustained maneuvers. As the maneuver length increases, the thrust goes to afterburner as quickly as possible to keep the kinetic energy high. Even so, the drag, and hence bleed rate, is so high that achieving afterburner faster does the unlimited model little good. For short runs, the unlimited thrust history was bang-bang. The initial

spool to idle was an attempt to slow the aircraft to approach corner speed and to allow more angle of attack without violating load factor limits. The surge to afterburner was then an effort to lessen the bleed rate as maximum angle of attack was neared. The nominal thrust history remained nearly constant at an average of idle and afterburner thrust because the engine could not spool fast enough for the optimal bang-bang. The bang-bang strategy did not help much, though, because of the tremendous inertia present at the speeds involved

Roll and pitch agilities were much more important. Roll limits were most restrictive in lower hemisphere pointing where fast reorientation of the lift vector is required. To targets requiring a 90 deg bank, for example, the unlimited roll rate model was able to point about 10 deg farther in the same time regardless of total angle traveled. This is because the aircraft cannot start its pull-up until the bank is essentially completed.

Removing pitch rate limits produced much faster times only in trajectories of moderate length (total angle less than 100 deg), in which quick pop-ups to maximum angle of attack are not penalized by excessive speed loss. For example, in 4 s the unlimited pitch rate model travels about 10 deg farther than nominal, regardless of maneuver plane. Over longer maneuvers, energy-maneuverability necessitated conservative angles of attack to keep the airspeed bleed rate low. In these cases, the unlimited angle of attack remained lower than nominal until the end to avoid drag penalties. There was little advantage in the advanced pitch capability here.

Standard maneuvers may conceal an aircraft's true performance capability and should not be used for agility comparison. In the case of an unlimited pitch rate model performing standard maneuvers, the aircraft must roll at 1 g regardless of pitch rate limits. The extra time spent waiting for the roll to complete before loading the aircraft annuls any advantage from the enhanced pitch capability.

The research suggests some agility metrics for pointing maneuvers and discounts others. The major trend from the time histories was that models acquiring g_{max} sooner and holding it longer pointed faster. Thus time to acquire g_{max} and percentage of maneuver time spent there might be valid new parameters. The unlimited roll rate studies confirm that time to roll 180 deg is a good performance measure for some targets.

Maximum turn rate and minimum turn radius are not valid, however, for these very dynamic maneuvers. They occur at the corner speed, the speed at which the lift limit and g_{max} lines intersect in a doghouse plot. For the aerodynamic model this speed is about 810 ft/s. There is no indication from the time histories that the aircraft tries to reach and maintain 810 ft/s. It frequently flies far below this to 600 ft/s.

This research promotes optimal control as a way of studying agility. Constraints can easily be added or removed and any quantifiable parameter(s) can be optimized. The

most meaningful agility metrics can then present themselves through time histories and contour plots.

5.2 Recommendations

There are four major recommendations for further research:

1. Because the time spent at maximum load factor was found to be so important, pointing simulations for aircraft with different load factor limits should be performed.
2. Combat simulations involving the nominal model versus each of the unlimited models should be performed. These should affirm that the unlimited models do indeed achieve first firing opportunity because of their superior pointing capability. They would also reveal weaknesses in the unlimited models, particularly if the first shot were to miss and a sustained battle resulted. This is because the unlimited models tend to trade energy for position. Combat simulations would show if additional constraints such as minimum total energy or minimum airspeed needed to be placed on the nominal model.
3. Runs with a large number of nodes should be made to further quantify the effects of the cubic spline fit. Switching to an algorithm that does not require trajectory segmentation (such as successive sweep or differential dynamic programming) might also be tried.
4. The simulations were performed assuming angle of attack, bank angle, and thrust could be directly controlled, ignoring actuator and control surface dynamics. The next step is to design a control system taking these extra degrees of freedom into account.

References

1. Herbst, W. B., "Dynamics of Air Combat," Journal of Aircraft, Vol. 20, No. 7, July 1983.

2. Skow, Andrew M. and Hamilton, Bill, "Advanced Fighter Agility Metrics," *Proceedings of the 29th Society of Experimental Test Pilots Symposium*, Beverly Hills, CA, 1985.

3. Bitten, R., "Qualitative and Quantitative Comparison of Government and Industry Agility Metrics," AIAA 89-3389-CP, *Proceedings of the AIAA Atmospheric Flight Mechanics Conference*, Boston, MA, August 14-16, 1989.

4. Dorn, M. D., "Aircraft Agility: The Science and the Opportunities," AIAA 89-2915, *Proceedings of the AIAA, AHS, and ASEE Aircraft Design, Systems, and Operations Conference*, Seattle, WA, July 31-August 2, 1989.

5. Riley, David R. and Drajeske, Mark H., "An Experimental Investigation of Torsional Agility in Air-to-Air Combat," AIAA 89-3388-CP, *Proceedings of the AIAA Atmospheric Flight Mechanics Conference*, Boston, MA, August 14-16, 1989.

6. McAtee, Thomas P., "Agility in Demand," *Aerospace America*, Vol. 26, No. 5, May 1988.

7. Cannon, Lt. John, "X-29 Agility: An Optimal Approach," Presented at Agility WORKshop, Wright-Patterson AFB, August 1989.

8. Shaw, Robert L., *Fighter Combat Tactics and Maneuvering*, Naval Institute Press, Annapolis, Maryland, 1985.

9. Vinh, Nguyen X., *Optimal Trajectories in Atmospheric Flight*, Elsevier Scientific Publishing Company, Amsterdam, 1981.

10. Etkin, Bernard, *Dynamics of Flight - Stability and Control*, John Wiley & Sons, New York, 1982.

11. Roskam, J., *Airplane Flight Dynamics and Automatic Flight Controls*, Roskam Aviation and Engineering Corporation, Ottawa, KS, 1982.

12. Hargraves, Charles and Paris, Steve, "Direct Trajectory Optimization Using Nonlinear Programming and Collocation," *Journal of Guidance, Control, and Dynamics*, Vol. 10, No. 4, July-August 1987.

13. Paris, S. W. and Hargraves, C. R., "Optimal Trajectories by Implicit Simulation, Volume I - Formulation Manual," AFWAL-TR-88-3057, November 1988.

14. Gill, Philip E., et al, "User's Guide for NPSOL (Version 2.1): A Fortran Package for Nonlinear Programming," Technical Report SOL 84-7, Systems Optimization Laboratory, Stanford University, September 1984.

15. Paris, S. W., et al, "Optimal Trajectories by Implicit Simulation, Volume II - User's Manual," AFWAL-TR-88-3057, November 1988.

16. Gill, Philip E., Murray, Walter, and Wright, Margaret H., *Practical Optimization*, Academic Press, London, 1981.

17. Scales, L. E., *Introduction to Non-Linear Optimization*, Springer-Verlag New York Inc., New York, 1985.

18. Dyer, Peter and McReynolds, Stephen R., *The Computation and Theory of Optimal Control*, Academic Press, New York, 1970.

19. McAtee, T. P. , "Agility - Its Nature and Need in the 1990s," *Proceedings of the 31st Society of Experimental Test Pilots Symposium*, Beverly Hills, CA, September 23-26, 1987.

20. Dole, Charles E., *Flight Theory and Aerodynamics - A Practical Guide for Operational Safety*, John Wiley & Sons, New York, 1981.

21. DeMeis, Richard, "Fighters Get a Quick Draw," *Aerospace America*, Vol. 24, No. 9, September 1986.

22. Kalviste, J., "Point-and-Shoot Agility Parameter," Briefing Presented at the Workshop on Agility Metrics held at the AF Flight Test Center, Edwards AFB, CA, March 8-10, 1988.

23. Tamrat, B. F., "Fighter Aircraft Agility Assessment Concepts and Their Implication on Future Agile Fighter Design," AIAA-88-4400, *Proceedings of the AIAA/AHS/ASEE Aircraft Design, Systems, and Operations Conference*, Atlanta, GA, September 1988.

24. Scott, William B., "Air Force, NASA Conduct Tests to Define Fighter Aircraft Agility," *Aviation Week & Space Technology*, Vol. 130, January 9, 1989.

25. Lawless, Capt. Al, Personal Discussion, February 2, 1990.

26. Press, William H., et al, *Numerical Recipes - The Art of Scientific Computing*, Cambridge University Press, Cambridge, England, 1989.

27. Burgin, George H., "Improvements to the Adaptive Maneuvering Logic Program," NASA Contractor Report 3985, 1986.

28. Kerrebrock, Jack L., *Aircraft Engines and Gas Turbines*, The MIT Press, Cambridge, MA, 1981.

29. Churchill, Ruel V. and Brown, James Ward, *Fourier Series and Boundary Value Problems*, McGraw-Hill Book Company, New York, 1978.