

BAR-HALO FRICTION IN GALAXIES. II. METASTABILITY

J. A. SELLWOOD

Department of Physics and Astronomy, Rutgers University, 136 Frelinghuysen Road, Piscataway, NJ 08854-8019;
sellwood@physics.rutgers.edu

AND

VICTOR P. DEBATTISTA¹

Department of Astronomy, University of Washington, Box 351580, Seattle, WA 98195-1580; debattis@astro.washington.edu

Received 2005 June 5; accepted 2005 November 8

ABSTRACT

It is well established that strong bars rotating in dense halos generally slow down as they lose angular momentum to the halo through dynamical friction. Angular momentum exchanges between the bar and halo particles take place at resonances. While some particles gain and others lose, friction arises when there is an excess of gainers over losers. This imbalance results from the generally decreasing numbers of particles with increasing angular momentum, and friction can therefore be avoided if there is no gradient in the density of particles across the major resonances. Here we show that anomalously weak friction can occur for this reason if the pattern speed of the bar fluctuates upward. After such an event, the density of resonant halo particles has a local inflexion created by the earlier exchanges, and bar slowdown can be delayed for a long period; we describe this as a metastable state. We show that this behavior in purely collisionless N -body simulations is far more likely to occur in methods with adaptive resolution. We also show that the phenomenon could arise in nature, since bar-driven gas inflow could easily raise the bar pattern speed enough to reach the metastable state. Finally, we demonstrate that mild external or internal perturbations quickly restore the usual frictional drag, and it is unlikely therefore that a strong bar in a galaxy having a dense halo could rotate for a long period without friction.

Subject headings: galaxies: evolution — galaxies: formation — galaxies: halos —
galaxies: kinematics and dynamics — galaxies: spiral — methods: numerical

1. INTRODUCTION

In earlier work (Debattista & Sellwood 1998, 2000, hereafter DS98 and DS00, respectively), we showed that strong, rapidly rotating bars in dense halos slow down quickly due to dynamical friction. Our finding of strong friction is consistent with theoretical work (e.g., Weinberg 1985, 2004; Hernquist & Weinberg 1992), and the consequent braking of bars is reported in other fully self-consistent simulations (e.g., Little & Carlberg 1991; O’Neill & Dubinski 2003).

Most strong bars in real galaxies appear to rotate rapidly, in the sense that the dimensionless ratio $\mathcal{R} \lesssim 1.4$; here $\mathcal{R} \equiv R_c/a_b$, where R_c is the radius of corotation and a_b is the semimajor axis of the bar.² This ratio is not easy to determine directly from observation, but Aguerri et al. (2003) summarize results from a number of galaxies, and Debattista & Williams (2004) add a new result with much lower uncertainty obtained by an integral field method (but see also Rautiainen et al. 2005). Indirect evidence comes from the location of dust lanes (Athanasoula 1992) and rings (Buta & Combes 1996).

In DS98 and DS00, we combined our finding that bars in dense halos soon become slow with the observation that bars in real galaxies are fast to attempt to constrain the contribution to the central attraction that comes from the inner dark matter halo. But others have challenged the result that bars become slow, which calls the constraint on halo density into question. The criticism by

Athanasoula (2003) is more of interpretation than substance, since strong bars are subject to fierce braking in her simulations, and bars that experience little friction are weak. We will report \mathcal{R} values for models similar to hers in Paper III of this series (J. A. Sellwood & V. P. Debattista 2006, in preparation).

Valenzuela & Klypin (2003, hereafter VK03), on the other hand, claim counterexamples of strong bars in dense halos that stay fast for cosmologically interesting periods of time. VK03 argue that their different result, which disagrees with all previous simulations and with the theoretical work cited above, is in fact correct and suggest that only their simulations have the numerical resolution to reveal the proper behavior.

When started from their initial conditions, simulations with our code (§ 2) behave in many respects as reported by VK03; they form strong bars of similar lengths and pattern speeds, for example. However, the bars in our simulations generally exhibit strong friction and quickly become unacceptably slow. One of our experiments is anomalous, however, and shows an even longer delay in the onset of friction than found by VK03. Although our anomalous result may be artificial (it is not reproducible when numerical parameters are changed), it presented us with an opportunity to discover how friction can be avoided for long periods.

Tremaine & Weinberg (1984) showed that dynamical friction in a quasi-spherical system arises because of resonant interactions between a rotating potential perturbation and the orbits of particles. As the decreasing bar pattern speed sweeps across a resonance with some halo orbits, their angular momenta may be substantially changed (Sellwood 2006, hereafter Paper I). Halo particles may either gain or lose angular momentum as they cross a resonance, and to first order there is no net loss or gain. However, to second order in the perturbing potential there is usually a

¹ Brooks Fellow.

² It is the value of this ratio today that matters, and not the fact that the bar has slowed a lot, as was misstated by Athanasoula (2003). Finding evidence of the past history of the bar pattern speed would be an even greater observational challenge.

TABLE 1
NUMERICAL PARAMETERS FOR RUNS F AND S

Parameter	Cylindrical Grid	Spherical Grid
Grid size.....	$(N_R, N_\phi, N_z) = (81, 128, 125)$	$N_r = 300$
Angular components.....	$0 \leq m \leq 8$	$0 \leq l \leq 4$
Outer radius (kpc).....	29.4	350
z -spacing (pc).....	21.875	...
Softening length ϵ (pc).....	43.75	...
Number of particles.....	199,996	3,351,798
Shortest time step (Myr).....	0.14	0.14

net gain in angular momentum by the halo particles, leading to a friction-like drag on the bar even in a perfectly collisionless system. The bias arises because the number density of halo particles is usually a decreasing function of angular momentum, leading to the excess of gainers over losers.

The distribution of particles about the principal resonances responsible for the torque on the bar is altered by the evolution itself, as shown in Paper I. In particular, exchanges between the halo particles and the bar combine with the time-dependent pattern speed to cause the density of particles in the vicinity of the resonance to develop a pronounced shoulder, with the instantaneous resonance lying on the high angular momentum side of the shoulder, where the gradient is negative.

In § 3, we show that should Ω_p fluctuate upward for some reason after a period of normal friction, the resonance may cross to the other side of the previously created shoulder, where the gradient with angular momentum may be locally flat or even reversed. In these circumstances, net exchanges at the resonance will no longer lead to friction. Since friction is generally dominated by a few resonances, which all behave in a similar manner, such a change can lead to a dramatic decrease in the net torque.³

We show that an upward fluctuation in the bar pattern speed in collisionless N -body simulations is not reproducible when numerical parameters are changed. The pattern speed fluctuates upward before periods of weak friction in experiments A₁ and B reported by VK03. In § 4 we show that adaptive mesh refinement is the likely culprit for a numerical artifact causing their anomalous results.

The possibility that bars in real galaxies with dense halos could experience little friction would be of great interest if there were a physically realistic reason for the bar pattern speed to rise. We show in § 5 that gas inflow in bars may have such a consequence, but we also find (§ 6) that other physically relevant factors are likely to prevent the low-friction state from persisting for long.

We describe the low-friction state as metastable both because it relies on a local feature in the phase space density and because it is fragile. It should be noted that metastability applies only to anomalously weak friction on a strong bar in a dense halo—friction will always be mild when the bar is weak or the halo density low.

2. SIMULATIONS

Model A₁ presented by VK03 had an exponential disk embedded in a cosmologically motivated, cold dark matter halo that had an approximately NFW (Navarro, Frenk, and White) density profile (Navarro et al. 1997) with a concentration index $c \approx 15$

after compression. These authors kindly made available the initial positions, velocities, and masses of all the particles in their model, and we have run a large number of simulations from these initial conditions using our hybrid, polar-grid code (Sellwood 2003). The numerical parameters used for two runs (runs F and S explained below) are given in Table 1, and we varied the numerical parameters in other runs. We employ the units used by VK03; the exponential disk has a mass of $4.28 \times 10^{10} M_\odot$ and scale length of 3.5 kpc; the halo has a mass of $2.00 \times 10^{12} M_\odot$.

The evolution of the pattern speed, Ω_p , in many of our simulations using their initial particles is shown in Figure 1; the dot-dashed curve shows the result reported by VK03 for this model using their ART code (Kravtsov et al. 1997). In contrast to their result, the bars in all but one of our cases slow quite quickly, reaching $\sim 16 \text{ km s}^{-1} \text{ kpc}^{-1}$ by ~ 5 Gyr, whereas VK03 found that the pattern speed remained roughly constant, $\Omega_p \simeq 27 \text{ km}^{-1} \text{ s}^{-1} \text{ kpc}^{-1}$, until ~ 4 Gyr and then decreased below $20 \text{ km s}^{-1} \text{ kpc}^{-1}$ by ~ 6.5 Gyr. In fact, the slowdown *rate* in all experiments is remarkably similar, the different curves being approximately parallel, with a generally small time offset. Apart from the delay, the decrease reported by VK03 is roughly consistent with the drop we observe in most cases.

However, in one of our experiments, which we denote run F (for “fast”), Ω_p did not decrease significantly until almost 12 Gyr from the start, although when it did start to decrease, the rate of decline was again quite comparable to that in the majority of our experiments and that found by VK03. As already noted, our experiments differ from each other only by the numerical parameters adopted. The parameters for run F (Table 1) are typical of all the experiments shown, and the different behavior is not a question of a lack of numerical convergence; parameters (time step, grid spacing, softening, etc.) in the other experiments were both refined and made coarser, from the set that led to the anomalous result.

In all but one case, we used all the particles supplied by VK03; in the remaining case (Fig. 1, *dotted line*), we employed the full number of disk particles but only every 10th halo particle, which we made 10 times more massive. The result, even with this quite

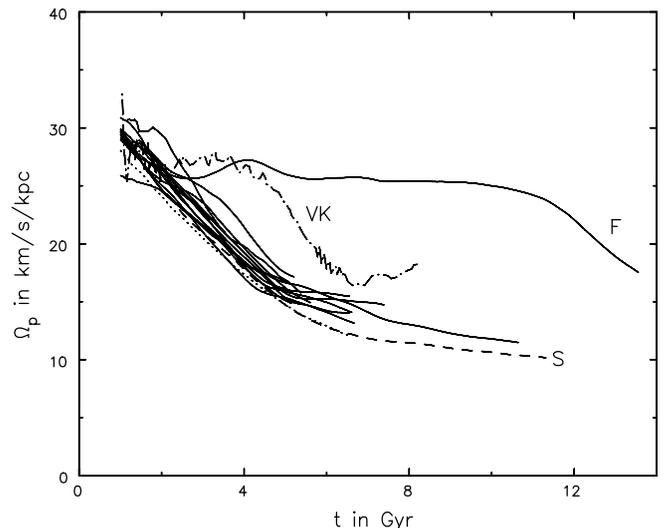


FIG. 1.—Pattern speed evolution in many experiments using our code that were all begun from the initial model A₁ of VK03. Each line shows a separate simulation that differs from the others only in the numerical parameters adopted. Curves from our runs F and S are so labeled, and the dot-dashed line labeled VK shows the evolution, reproduced from Fig. 10 of VK03, when this model was simulated with their ART code.

³ Holley-Bockelmann & Weinberg (2005) construct a halo with no gradient at the outset and report mild friction on the bar, but such models are quite contrived.

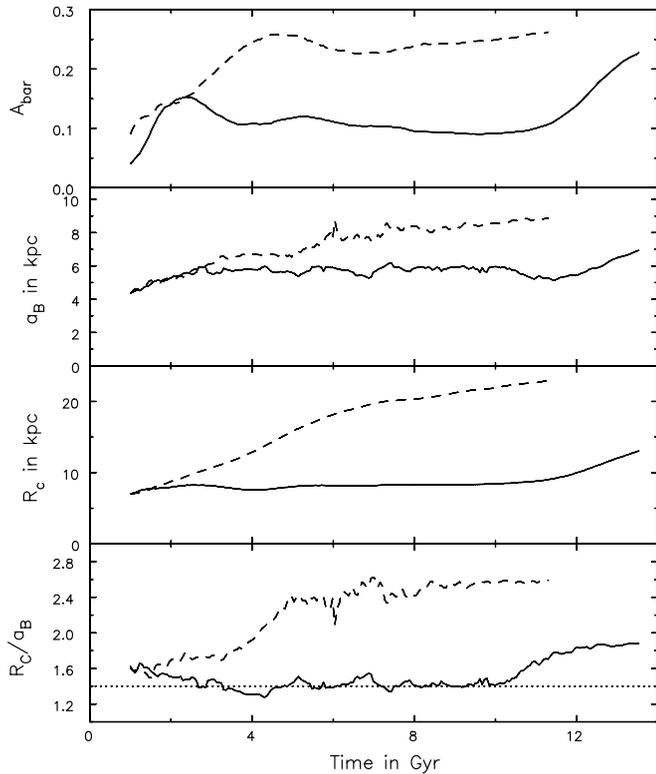


FIG. 2.—*Top to bottom*: Bar amplitude, the bar length, radius of corotation, and the ratio \mathcal{R} . The solid curves show these quantities for run F, while the dashed curves show them for run S. The horizontal dotted line in the bottom is drawn at the value $\mathcal{R} = 1.4$.

drastic reduction in numerical quality, is that the bar slowed in a similar manner as shown.

A further experiment, shown by the dashed curve in Figure 1, has the same numerical parameters as run F. In this experiment, we rotated the radius vector to each disk particle through a random angle before we began, effectively making a different draw of disk particles from the parent population, which leads to a different spectrum of initial disk density fluctuations with a similar amplitude. In this case, even with the same numerical parameters as in run F, the bar slows, in general agreement with the other results. We denote this experiment run S (for “slow”).

Aside from the anomalous run F, the simulations shown by solid lines in Figure 1 have numerical parameters that were changed from those listed in Table 1 for runs F and S. The differences were as follows. We both doubled and halved the basic time step, we increased the resolution of the polar grid, reduced the spacing of the grid planes, halved the softening parameter (on the finer grid), both increased and decreased the number of shells in the spherical grid, and both reduced and increased l_{\max} to 2 and 8 on the spherical grid. Finally, we imposed reflection symmetry about the disk midplane in one of these runs. In all these runs, Ω_p declined at approximately the same rate as in run S.

The solid curves in Figure 2 show the time evolution of the bar amplitude, bar semimajor axis a_b , corotation radius R_c , and the ratio $\mathcal{R} = R_c/a_b$ for run F. (Our procedure for making these measurements from simulations is described in the Appendix.) The dashed curves, on the other hand, show the same quantities from run S; it is clear that the bar in run S soon becomes, and remains, unacceptably slow. We find that all other models from Figure 1 that slow early show a rapid rise in \mathcal{R} similar to that in run S.

As reported by VK03 for their bar, the bar in our run F remains acceptably fast, $\mathcal{R} \simeq 1.4$, for a long time, although \mathcal{R} increases

when the bar finally begins to slow. It can be deduced from Figures 10 and 14 of VK03 that the corotation radius has increased from ~ 6.9 to ~ 10.5 kpc, while they report (their § 7) that the bar length in their model A₁ is 6–6.5 kpc after 8.5 Gyr of evolution. Thus, the final value of \mathcal{R} in their simulation is indeed significantly larger than 1.4.

Thus, our experiments shown in Figure 1 are in excellent agreement with almost all aspects of the results obtained by VK03 using their entirely different numerical technique, except for the time at which Ω_p decreases from ~ 27 to < 20 km s⁻¹ kpc⁻¹. In most of our experiments, this happened earlier than found by VK03, but it occurred later in our anomalous run F. We even concur that the bar’s corotation circle remains acceptably close to the bar end in run F for as long as Ω_p remains ≥ 25 km s⁻¹ kpc⁻¹. We account for our discrepant result in § 3 and that obtained by VK03 in § 4.

3. METASTABILITY

3.1. Restricted Experiment

Lin & Tremaine (1983) used restricted N -body simulations to show that driving the system with a perturbation (a companion galaxy in their case) at a constant frequency considerably delayed the onset of friction once unforced evolution was allowed. They suspected, but were unable to show, that the absence of friction was due to all the resonant particles having been scattered. Here we show friction can also be suppressed in the bar case. At first, we again suspected that the absence of friction was due to cleared resonances, but this turned out not to be the case, and we demonstrate that it is due to an adverse gradient of particle density about the major resonances.

Figure 3 shows the evolution of Ω_p of a rigid bar in a test particle halo. The halo, represented by 10^7 particles, has a Hernquist (1990) density profile with an isotropic distribution function, and the uniform-density, ellipsoidal bar has a mass that is 2% of the halo mass and an axis ratio 1 : 0.2 : 0.05, with the long axis equal to the scale radius of the Hernquist profile. (This is the fiducial experiment described in § 6 of Paper I, where further details are given.)

The dotted curve shows the evolution of Ω_p dictated by conservation of angular momentum for a fixed moment of inertia, as the bar experiences dynamical friction. The solid curve shows what happens when we drive Ω_p back up to 0.4 between times 180 and 200. The bar pattern speed, which is free to evolve as a result of halo friction after $t = 200$, stays approximately constant

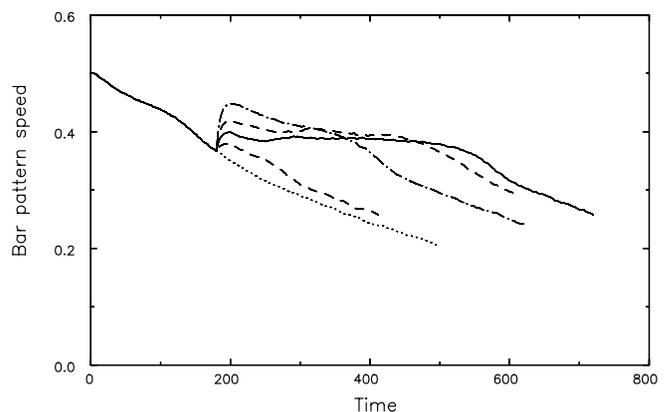


FIG. 3.—Time evolution of Ω_p of a rigid bar in an experiment with non-interacting halo particles. The dotted curve shows the unforced evolution. The bar pattern speed is driven back up between times 180 and 200 to 0.4 (solid line), to 0.38 and 0.42 (dashed lines), and to 0.45 (dot-dashed line).

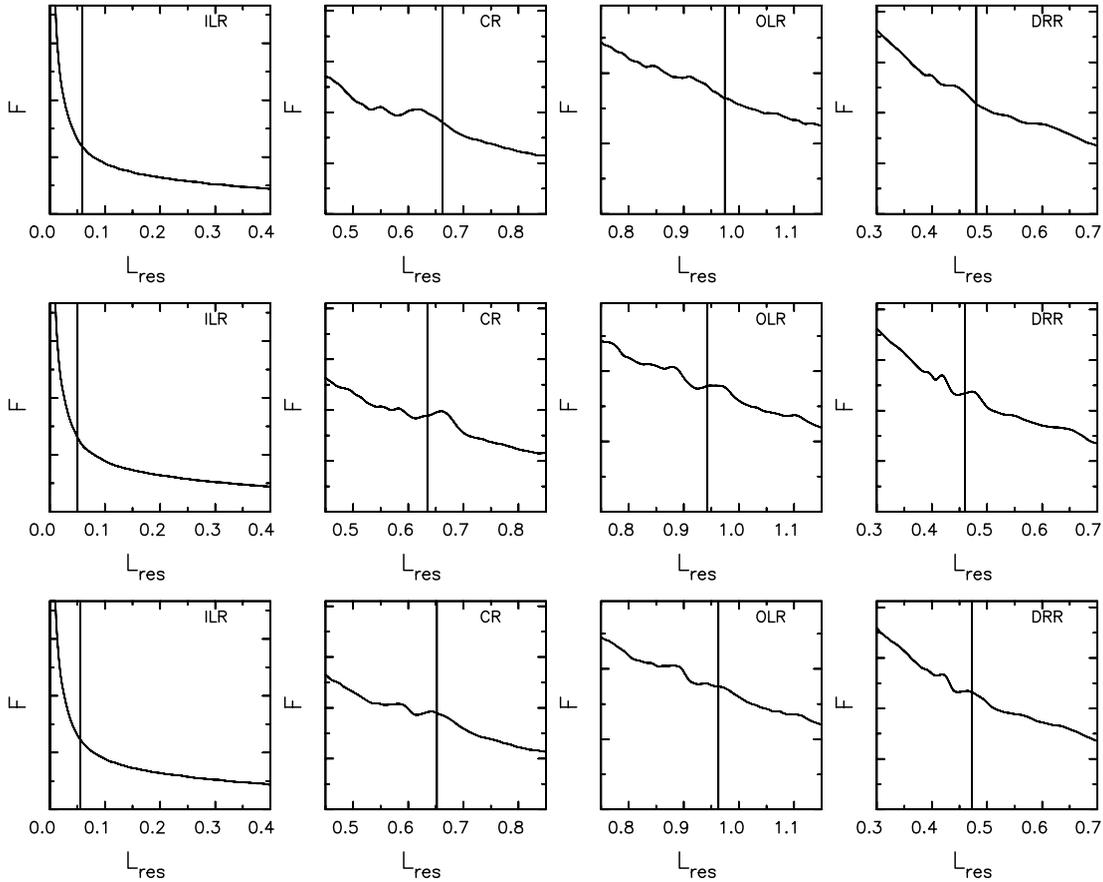


FIG. 4.—Mean density F of particles as a function of the resonant angular momentum, L_{res} , at the four major resonances at three different times in the model shown by the solid line in Fig. 3. The columns of panels are labeled by the resonances; the top row is for time 176, during the period of normal friction, the middle row is for time 224, when friction is very weak, and the bottom row is at time 520, as normal friction starts to resume. The scaling of F is linear from zero.

for some time, and the onset of friction is considerably delayed. The other curves show the evolution when Ω_p is driven back up to different values.

3.2. Normal and Anomalous Gradients

In Paper I, we introduced the function $F(L_{\text{res}})$, which is the phase space density reduced to a function of a single variable by averaging over all orbit phases and eccentricities at fixed L_{res} . For any orbit, the frequency difference from resonance is $\Omega_s = (n\Omega_\phi + k\Omega_r)/m - \Omega_p$, where Ω_ϕ and Ω_r are the angular frequencies of the orbit in the unperturbed spherical potential (§ 3.1 of Binney & Tremaine 1987) and n , k , and m are integers. The quantity L_{res} is the angular momentum of a circular orbit that is the same distance in frequency, Ω_s , from the resonance as an orbit of arbitrary eccentricity—see Paper I for a fuller description. The four resonances that are most important for friction are the corotation resonance (CR), where $n = m$ and $k = 0$; the inner and outer Lindblad resonances (ILR and OLR), where $n = m = 2$ and $k = \mp 1$; and the direct radial resonance (DRR), where $n = 0$, $m = 2$, and $k = 1$.

Figure 4 shows the function $F(L_{\text{res}})$ near each of the four major resonances at three different times in the simulation shown by the solid line in Figure 3. As shown in Paper I, the distribution of particles about each major resonance is a generally decreasing function of L_{res} , but with a shoulder associated with those resonances that contribute strongly to friction. The top row is for $t = 176$, during normal friction, the middle row is at $t = 224$, during the metastable phase, and the bottom row is at $t = 520$, as normal friction is about to resume. We first discuss the situation

at corotation (CR; Fig. 4, *second column*). At $t = 176$, the distribution has a local maximum on the low- L_{res} side of the resonance, as was already reported in Paper I. When the pattern speed is driven back up, the resonance moves to the other side of the local maximum, where the gradient is slightly positive, as shown at $t = 224$; finally, as normal friction resumes ($t = 520$), the resonance is just passing the local maximum. Similar behavior can also be discerned at the OLR and DRR.

These figures make it clear that the rise in Ω_p was enough for dF/dL_{res} to have changed sign at the three resonances that are most important for friction. Friction is greatly reduced when the slope of $F(L_{\text{res}})$ becomes positive in the immediate vicinity of the resonances that dominate the angular momentum exchanges with the bar. The change in gradient removes the usual excess of gainers over losers that is responsible for friction.⁴ Strong friction does not resume until the gradients at the dominant resonances become decisively negative once more.

3.3. Long-Term Evolution

The solid line in Figure 3 is not precisely flat after $t = 200$, and a slow decrease in Ω_p is discernible. The mild friction that causes the slow decrease probably results from exchanges at higher order resonances. The set of possible resonances between a halo orbit and a rotating perturber is large, since a resonance arises for

⁴ The gradient of the usual distribution function, $\partial f/\partial L$ at constant radial action, must be nearly flat at the most important resonances in order that exchanges with the perturbing potential are neutral. The positive gradient in F is probably a result of its complicated relation to f .

any set of integers k , m , and n . Thus, even if the important low-order resonances cause no friction, some residual friction is provided by exchanges at the many higher order resonances; gradients in $F(L_{\text{res}})$ at these weak resonances were not much affected by the previous evolution. The weaker coupling between the perturbation and particles at higher order resonances (Hernquist & Weinberg 1992; Paper I) leads to slower, but nonzero, angular momentum transfer. (We have been unable to find convincing evidence of angular momentum exchanges at a number of possible resonances, which is hardly surprising, as friction is so weak.) Thus, the pattern speed gradually decreases until the full friction force picks up again when gradients of $F(L_{\text{res}})$ at the dominant resonances are negative once more.

The lower dashed line in Figure 3 shows friction was not greatly affected when the rise of Ω_p was insufficient to move the CR past the maximum of $F(L_{\text{res}})$. The upper dashed line, on the other hand, shows that friction reappears briefly after Ω_p is driven up to a slightly larger value, but again becomes anomalously weak for a long period. Analysis of $F(L_{\text{res}})$ in this case shows that the slight initial drop in pattern speed allows the resonances to reach the adverse gradient in $F(L_{\text{res}})$. The duration of the metastable phase is slightly shorter because the local maximum in $F(L_{\text{res}})$ is eroded somewhat by the evolution after the pattern speed is driven up. After the pattern speed is driven up to a still higher value, the bar slows continuously with only a slight reduction in friction as the bar passes through $\Omega_p \approx 0.38$; i.e., too large a rise in Ω_p does not lead to metastability.

3.4. Relevance to Figure 1

It therefore seems likely that the long period of steady bar rotation reported by VK03 and the even longer such period in our run F were caused by the upward fluctuations in Ω_p in both cases visible in Figure 1. These upward fluctuations must have raised Ω_p to values at which the gradients in $F(L_{\text{res}})$ at the dominant resonances are reversed.

Figure 5 shows this to be the case in our run F. In order to obtain this figure, we needed to estimate L_{res} for each particle in this fully self-consistent simulation with a disk and bar. Unlike in the restricted experiments just reported, the mass profiles of the disk and halo evolve continuously because a bar forms in the disk and the halo was not initially very close to equilibrium. We therefore approximated the potential well in the simulation as a time-averaged, spherical mass profile, extracted from the simulation over a short interval. We then computed the two frequencies of every halo particle from its energy and angular momentum defined by its instantaneous position and velocity at a moment in the middle of the period of the time averaging.

The features in Figure 5 are not quite as clear as those in Figure 4, possibly because we had to erase the nonspherical components of the disk, bar, and halo, in order to compute the orbit frequencies. (We had to discard a tiny fraction of particles that had less energy than that of a circular orbit of their angular momentum in this approximate potential.) Nevertheless, a feature that resembles the shoulder at corotation in Figure 4 can be seen, and the rise in pattern speed between the two times moves the resonance across the shoulder to change the gradient in $F(L_{\text{res}})$.

We can only speculate why there was a longer delay before full friction resumed in our run F than in the experiment reported by VK03. One possible reason might be that the upward fluctuation in their run may have taken Ω_p only just past the local maximum of $F(L_{\text{res}})$.

It is also possible that collisional relaxation in simulations with self-gravity allows halo orbits to diffuse slowly in phase space, thereby eroding the local maximum in $F(L_{\text{res}})$ (which could not

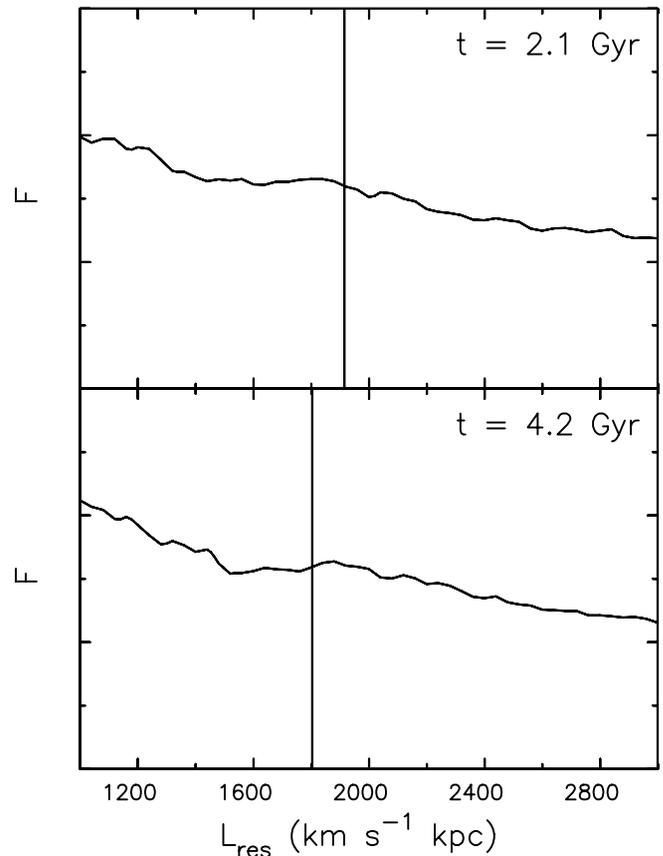


FIG. 5.—Function $F(L_{\text{res}})$ at corotation at two different times in run F from Fig. 1. The earlier time, $t = 2.1$, is during normal friction, while $t = 4.2$ is after the pattern speed has risen.

have happened in the restricted experiments of § 3.3). This effect seems to be minor in our run F, since the higher order resonance explanation accounts for the resumption of friction on the appropriate timescale (compare Figs. 1 and 3). In fact, if the worry over rapid orbit diffusion raised by Weinberg & Katz (2002, 2005) were important, the metastable state could not persist for long (see also § 7.4). Note that the absence of friction is not simply a question of inadequate coverage of phase space by the finite number of particles, since the bar still slows as normal (Fig. 1, dotted curve) when the density of particles is reduced by a factor of 10 (see also Paper I).

3.5. What Happened in Run F?

We have attempted to understand what caused the upward fluctuation in our run F. Detailed comparison of this run with others that slowed has led us to conclude that it was caused by a random interaction of the bar and a spiral pattern in the disk that occurred at a rare relative phase so as to add angular momentum to the bar—spirals generally have the opposite effect at most relative phases.

It should be noted that most spiral activity in disk simulations is genuinely stochastic. As the initial spiral patterns are determined by the swing-amplified (Toomre 1981) seed spectrum of particle noise laid down by the random selection of particles, the very first features are insensitive to parameters such as grid resolution. However, the subsequent patterns can differ macroscopically as the result of seemingly insignificant changes to numerical parameters. Slight changes to the disk responsiveness or central attraction cause small differences to the early evolution,

but the subsequent sequence of spiral patterns in models that differ in this way visibly diverges in remarkably few dynamical times.

These differences are largely responsible for the scatter in slowdown times (but not rates) already noted in Figure 1. Fortunately, the statistical properties of the disk at later times are not sensitive to these differences, and the divergent sequences of spirals result in similar disk random motion and—excepting run F—bar speed evolution.

4. ADAPTIVE REFINEMENT

The evidence presented so far indicates that stochastic disk evolution can occasionally, with our code, put the bar into the metastable state, which was also reached in the simulation reported by VK03. It is possible they were just unlucky to find this state, although they report little bar friction over the (short) period of evolution in their model B also, making the chance explanation seem unlikely.

The principal numerical difference between their method and ours is their use of adaptive mesh refinement (Kravtsov et al. 1997), and we now show that this feature may make the metastable state easier to reach. Our reasoning is as follows. As a bar develops, particles become more closely packed in the bar than they were in the disk, and an adaptive code will cause short-range gravitational forces in the bar to strengthen over those pertaining if resolution/softening were fixed.⁵ Stronger forces, or steeper potential gradients toward the center, have two effects: they cause the bar to contract, and they cause the orbital periods of particles at the same mean galactocentric distance to be somewhat shorter. Both effects will raise the pattern speed of the bar above what it would have been in a fixed resolution code, which could possibly be sufficient to push the bar into the metastable state.

4.1. Mimicry

To test this hypothesis, we ran a suite of further experiments with model A₁ from VK03, using different softening lengths, ϵ , and higher grid resolution.⁶ The bar slows normally in the basic model of this suite, shown by the dashed line in Figure 6. The softening length in this case, $\epsilon = 0.05R_d = 175$ pc, is 4 times greater than for most of those in Figure 1 and for the run shown by the dot-dashed curve in this Figure 6 that also slows normally.

But the bar stayed fast, at least for a while, in the four cases shown by solid lines when softening was reduced by a factor of 4 after different periods of evolution with the larger ϵ . In each of these four cases, softening was decreased abruptly at one instant in the interval $2 \text{ Gyr} \lesssim t \lesssim 2.8 \text{ Gyr}$, i.e., during the later period of bar growth. (In one of these four cases, the bar speed experiences a second upward fluctuation and then declines; we comment on that case in § 6.) The curve shown by the dotted line is for another case in which softening was merely halved—again the bar stayed fast.

This crude stratagem is intended to mimic the effects of adaptive refinement, although in our code the change is uniform over the inner cylindrical grid and abrupt. The change in softening affects only the small fraction of disk particles in the inner grid; the polar grid has cell sizes that increase with cylindrical

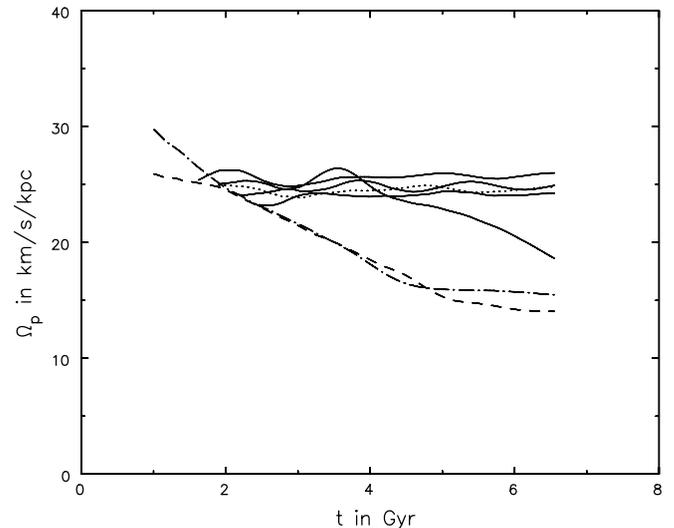


FIG. 6.—Time evolution of Ω_p in a fully self-consistent simulation with 4 times larger ϵ and finer grid than were used for run F (dashed line). The dot-dashed line shows the evolution with the same softening length as run F, but run on the finer grid. The solid lines show results when ϵ is reduced to the lower value after some evolution with the larger softening. The dotted curve shows the effect of only halving ϵ .

radius and interparticle forces are soon limited more by the grid than by the formal softening. It should be emphasized that this change does not affect forces on halo particles, which are computed from the spherical grid. As a result, the abrupt change in softening has a barely noticeable effect on the equilibrium of the model.

The surface density along the ridge of the inner part of the bar is almost 4 times higher than the original axisymmetric disk (Fig. 15 of VK03) well after the bar has formed. (We find a somewhat larger initial increase that relaxes back later as the bar becomes slightly rounder.) Since this surface density change takes place before the bar thickens through buckling instabilities, the mean distance between disk particles in this region decreases by about a factor 2 as the bar forms. Thus, a reduction of softening by a factor of 2 in our code is approximately equivalent to adaptive refinement in the ART code employed by VK03 for the same model. The effect of the change in softening strongly suggests that the metastable state that appears in several of the experiments reported by VK03 is a numerical artifact of adaptive refinement. Their finding that bars slow as normal if they limit the level of mesh refinement (their § 7) reinforces this suspicion.

VK03 report an additional model, their model C, which had more particles, and the highest resolution of their adaptive scheme was limited to 100 pc, as opposed to 44 pc for their other simulations. It is therefore unlikely that the pattern speed in this simulation would rise due the numerical artifact that affected model A₁ and possibly also their model B. Unfortunately, they did not run their simulation of model C for very long; the bar pattern speed is declining over the second half of the evolution, and the bar would probably have become slow had they continued the calculation.

4.2. Discussion of Softening

Adjusting the softening parameter, or the grid resolution, dynamically clearly changes the behavior. Why should this be? Softening can be thought of as an operation first to convolve the particle distribution with a smoothing kernel and then to solve for the full Newtonian field of the smoothed density distribution. Ideally, the softening length, ϵ , should be set large enough that

⁵ The forces between particles in grid codes have a softened form inside ~ 3 mesh spaces (see, e.g., Fig. 14 of Sellwood & Merritt 1994).

⁶ Softening length and grid resolution can be varied independently in the convolution method used for our cylindrical polar grid. Changes to ϵ have almost no effect, however, unless the grid cell dimensions are $\lesssim \epsilon$, which, for a polar grid, is possible in the inner parts only. The softening length, ϵ , is larger than the (R, ϕ) cell dimensions out to $R \simeq 0.3R_d$ for the shortest value of ϵ used in Fig. 6, and this radius increases nearly linearly with ϵ for a fixed grid size.

many particles lie within one softening volume, but short enough so as not to smooth, or bias, gravitational potential gradients.

If ϵ is large compared to the mean separation of particles in a region of quasi-uniform density, then small changes in the width of the softening kernel will not alter the smoothed density, and the global gravitational potential will be unaffected. Conceptually, we could imagine increasing N as ϵ is reduced all the way to the double limit $N \rightarrow \infty$ as $\epsilon \rightarrow 0$ without affecting the potential at all, but only if ϵ starts small enough not to smooth any density gradient significantly. Thus, wherever the twin ideals of many particles per softening volume and a softening length shorter than any density gradient hold, adaptive refinement would not change the dynamics.

Unfortunately, these ideals are generally not achievable in practical simulations of realistic stellar systems. The desire to resolve steep density gradients causes simulators to prefer values of ϵ on the low side in order to maximize spatial resolution or to minimize bias, which leads to few particles per softening volume; for example, VK03 continue to subdivide any cell that contains more than four particles. The penalty for this preference is an increase in the variance in the gravitational potential, but the increase in relaxation rate, through the change to the Coulomb logarithm, is quite small.

However, global gravitational potential gradients in the disk plane depend on softening unless $\epsilon \ll z_0$, the vertical density scale height. Density gradients in the x - and y -directions are quite shallow, even in a bar, but vertical gradients in thin disks are a different matter. It is hard to ensure $\epsilon \ll z_0$ for a thin disk in any code.

The consequences of inadequate vertical resolution are twofold. First, weakened vertical forces increase the vertical oscillation period of disk particles. Such a bias has a small impact on the in-plane motion of particles, which are mostly decoupled from their vertical motions, *except* at later times, when the bar buckles. Second, in-plane potential gradients are weakened by softening (see, e.g., problem 6-5 of Binney & Tremaine 1987) unless $\epsilon \ll z_0$. Thus, the dynamics of bar formation is different in simulations with different, but fixed, softening lengths.

Initially, $z_0 \approx 140$ pc in model A₁ (not 250 pc as stated in Table 1 of VK03), and the disk does not thicken much until some time after the bar has formed. The effect of using a different, but unchanging, softening length in our code is that the bar forms a little earlier and with a higher Ω_p when the softening length is shorter. Steady bar friction results in both cases shown in Figure 6, braking the bars at similar rates. This mild dependence on softening length could be avoided only if $\epsilon \ll z_0$ and N were increased appropriately—an ideal that is numerically expensive for a thin disk.

The ART code used by VK03 for this model refined cells to a size of 22 pc, or an effective minimum softening of ~ 40 pc, but it was probably larger over most of the disk. Adaptive refinement must therefore increase the in-plane forces as the bar strengthens, causing the bar to speed up purely for this numerical reason. Thus, adaptive refinement may trigger the metastable state.

5. METASTABILITY IN REAL GALAXIES?

We have shown that long periods without friction can occur in purely collisionless N -body simulations if the bar pattern speed rises for some reason. This can happen occasionally at fixed softening, or spatial resolution, but it is more likely to happen when spatial resolution is adaptive; in any case, the expected braking is recovered when the numerical procedure or parameters are changed.

It is interesting to ask whether real bars could avoid friction in a similar manner. It seems possible that random spiral events, or

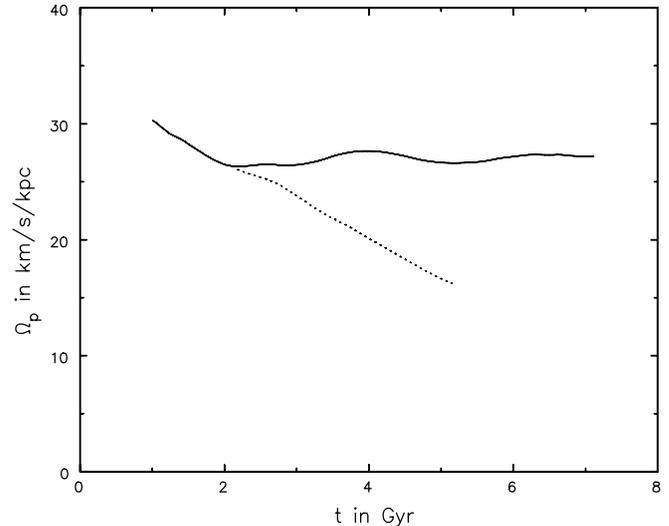


FIG. 7.—Time evolution of Ω_p in a fully self-consistent simulation (*solid line*) in which we increased the masses of the 1% most bound particles by 50% at $t = 2.24$ Gyr. The dotted curve shows the result without this change.

even the interaction with a minor satellite galaxy, could occasionally increase Ω_p and allow strong braking to be avoided for a while; but such behavior seems likely to be the exception, rather than the rule. However, real galaxies contain gas, which behaves differently from stars; gas is widely recognized to flow in toward the center of a bar. Large gas accumulations are found in regions a few hundred parsecs across (e.g., Sakamoto et al. 1999), which are sometimes resolved as inner gas rings. While the gas accumulations are a small fraction of the total disk mass, it is possible the mass increase in the center could well be enough to raise Ω_p .

To test this, we took a model in which the bar was slowing normally and increased the masses of each of the 1% most bound disk particles by 50% at $t = 2.24$ Gyr; i.e., we increased the total disk mass by 0.5% by adding mass at the bar center only. Figure 7 shows that frictional braking ceased and Ω_p stopped declining.

The central mass added in this experiment, $\sim 2 \times 10^8 M_\odot$, is quite consistent with observed gas masses in bar centers. It therefore seems entirely possible that gas inflow in bars could cause a bar to speed up enough to turn off dynamical friction, even in a dense halo.

6. FRAGILITY OF THE METASTABLE STATE

If the bar in an isolated galaxy with a dense halo has been spun up so that it experiences little friction, our simulations suggest that it might be several Gyr before the higher order resonances can reduce the pattern speed to the point at which strong friction resumes. However, real galaxies are not isolated and are subject to numerous small perturbations from infalling dwarf galaxies, as well as the passage of larger galaxies at greater distance; e.g., the Milky Way today has the Sagittarius dwarf, the Magellanic Clouds, and other dwarf companions. In addition, the Λ CDM model of galaxy formation predicts a much larger number of dark minihalos (Moore et al. 1998; Klypin et al. 1999) orbiting within the halo of a large galaxy. Since the absence of friction depends on a prearranged distribution of particles at the major resonances, the constant stirring of a galaxy halo by these perturbations might well cause strong friction to resume much earlier. We have conducted a number of experiments in order to investigate this possibility.

We perturbed our run F at $t = 4.2$ Gyr, which is in the metastable state, by a small satellite galaxy that flies by. We model the

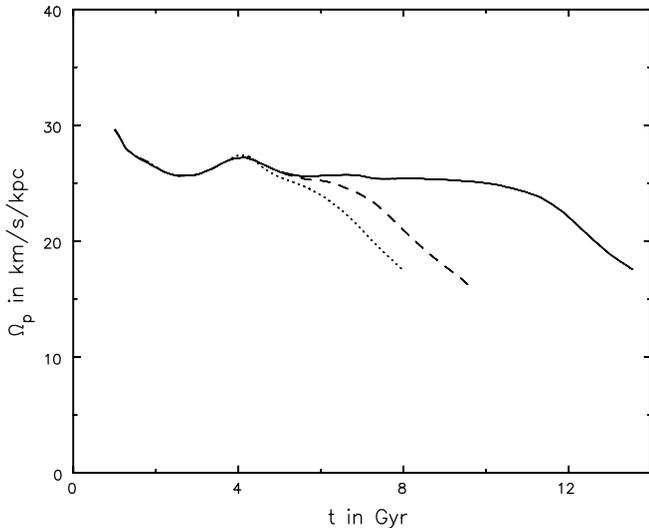


FIG. 8.—Time evolution of Ω_p in run F (*solid line*), again reproduced from Fig. 1. The dashed curve shows Ω_p when a 1% mass satellite galaxy flies by, passing a perigalactic distance of 30 kpc at $t = 4.2$ Gyr. The bar is jolted out of the metastable state and begins to slow. The dotted curve shows the bar slows after the disk is discontinuously rotated forward relative to the halo by 10° at $t = 4.2$ Gyr.

satellite as a rigid Plummer sphere with a scale radius equal to the exponential disk scale. We introduce it at a distance of 60 kpc from the center of the galaxy, at $t = 4.2$ Gyr, when its mass is very small, and increase its mass gradually to its final value at $t = 4.5$ Gyr. It has a mildly hyperbolic, polar orbit that crosses the disk plane at a radius of some 30 kpc at about $t = 4.9$ Gyr. In order to ensure that the main galaxy remains optimally centered in the grid, we shift the grid center every 16 time steps to coincide with the densest point (a feature that was enabled in the majority of the simulations already reported); all three components of total linear momentum are well conserved after the intruder has reached full mass. We have tried satellite masses that are 1%, 2%, and 5% of the total (disk + halo) mass of the main galaxy. Even for the 1% mass satellite, the halo is sufficiently perturbed for bar friction to resume, as shown by the dashed curve in Figure 8; the solid line again shows the evolution of Ω_p in run F.

In a separate experiment, we perturbed run F at the same moment by rotating every disk particle forward through 10° and then allowed the system to evolve freely once more. The purpose of this exercise is instantly to change the orientation of the bar relative to the density response in the halo, and therefore to change the orbital phases of the halo particles relative to the bar. While artificial, the sudden change of bar phase could perhaps resemble the effect of interaction with a strong spiral pattern.⁷ Just such an event seemed to be responsible for the upward fluctuation followed by early slowdown of one of the bars in Figure 6. The dotted curve in Figure 8 shows the evolution of this perturbed model; the upward fluctuation in Ω_p immediately after the imposed change is an artifact caused by our measuring the bar pattern speed from the slope of its phase angle with time, which includes our imposed 10° discontinuous change. Once this feature is outside our fitting window, Ω_p decreases at the usual rate. Thus, jolting the system in this manner again tipped it out of the metastable state.

⁷ The outer disk in our simulation supports generally rather weak spiral patterns after the bar has formed. Barred galaxies having gas and ongoing star formation would be expected to have stronger spirals.

The results of both experiments indicate that extremely mild disturbances to the system are sufficient to cause friction to resume. We conclude that the metastable state is highly fragile, and it is unlikely that it could persist in real galaxies for long.

7. FURTHER NUMERICAL ISSUES

7.1. Spatial Resolution

The spatial resolution of the fixed Cartesian grid used in our earlier experiments (DS98; DS00) was indeed quite low, and VK03 suggest that it was inadequate to capture the correct physics. We disagree.

As stated by Hernquist & Weinberg (1992) and shown in Paper I, dynamical friction is dominated by the quadrupole field of the bar, which couples most strongly to the low-order resonances. A good approximation to the correct quadrupole field can be obtained with a mesh of quite low spatial resolution, and the torque between a given bar and halo should be little affected by spatial resolution.

We have already presented supporting evidence for this statement, since we have shown that the rate of slowdown of the bar is insensitive to softening length and grid resolution. The different runs shown in Figure 1 have a range of softening lengths and grid resolutions, and ϵ differs by a factor 4 between the two runs with unchanging ϵ shown in Figure 6, which track each other remarkably closely. Furthermore, once friction picks up in the simulation by VK03 (Fig. 1; *dot-dashed curve*), the rate of slowdown is quite comparable to that we find in our simulations.

7.2. Bar Size

The pattern speed and spatial scale of the bar instability, and the corresponding properties of the resulting bar, depend on many factors. Two of the most important are the steepness of the inner rotation curve and the velocity dispersion of the disk.

VK03 suggest that inadequate spatial resolution can lead to longer bars. Indeed, if the central density gradient (of disk, bulge, and halo particles) is steep, inadequate spatial resolution will weaken the central forces, and this numerical bias leads to a rotation curve that rises more slowly than it should. Since more slowly rising rotation curves lead to longer bars (Sellwood 1981) with lower pattern speeds, the size of the bar may be artificially enhanced if inadequate spatial resolution smooths the sharper forces expected from a steep density gradient.

However, smaller bars are not the inevitable consequence of improved spatial resolution. The large size of the bars in DS00 was a consequence of the shallow density gradients in our models, which we deliberately selected because we were fully aware that the code we were using was unable to resolve steep gradients. Figure 9 shows the rotation curve and bar evolution of a model deliberately set up to have a halo with a large, low-density core to resemble the maximum-disk model reported in DS00, but evolved with our newer high-resolution code. Like the similar model in that paper, the disk has a Kuzmin-Toomre (KT) density profile with scale radius R_s , thickness $0.05R_s$, and initial velocity dispersion to make $Q = 1.0$ at all radii; the low-concentration halo has a polytropic density function and a mass about 5 times the disk mass and extends to $r = 28R_s$. The bar that forms in our high-resolution, hybrid-grid code is similar in size and amplitude to that reported in DS00 for the equivalent model run on their coarse grid, as shown in Figure 9 (*bottom*). Thus, bar size need not be related to spatial resolution.

DS00 compared the strengths and sizes of the bars in their models with bar properties determined from NGC 936. They concluded that the bars in their maximum-disk models were quite

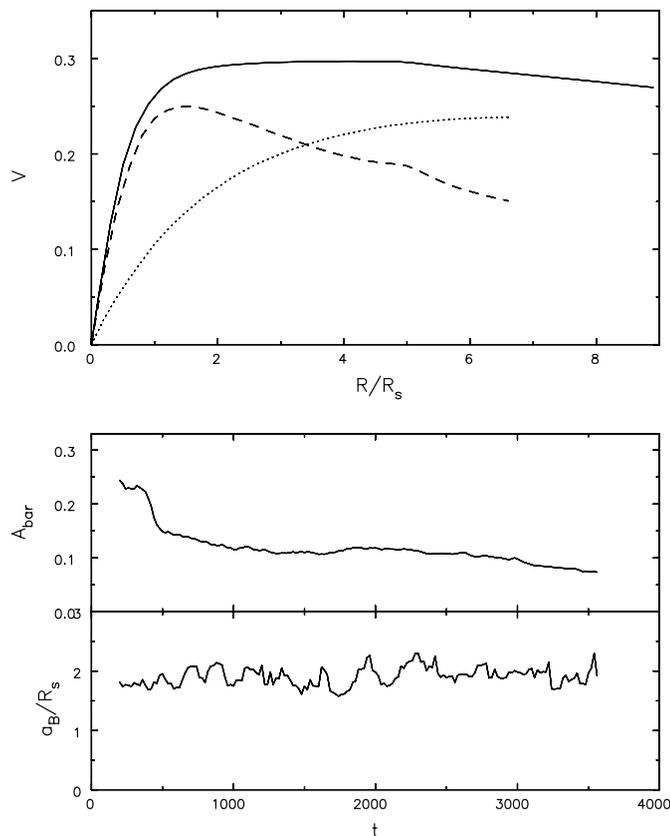


FIG. 9.—*Top*: Rotation curve (*solid line*) and separate contributions of disk (*dashed line*) and halo (*dotted line*) for a maximum-disk model, similar to that reported in DS00. *Bottom*: Bar amplitude and length as a function of time. Note that this initially large bar grows neither in length nor amplitude.

comparable to strong bars in real galaxies. Note also that the bar length in Figure 9 is reckoned relative to the length scale, R_s , of the initial KT disk; after the bar has formed, the outer disk is characterized by an exponential profile with a scale length about $2.5R_s$. This important difference has already been noted by VK03.

However, bars become longer and stronger when subject to strong frictional braking, as shown in Figure 2 and previously by DS00 and Athanassoula (2002). DS00 showed that bars that had been braked in denser halos were significantly stronger than that in NGC 936, which might be typical of strongly barred galaxies. Even the mild braking in their maximum-disk model strengthened the bar to the point where it was only marginally consistent with the data on NGC 936. We therefore agree with VK03 that bars often become quite unrealistically large and strong. The more realistic bars in models with maximum disks provide further evidence against the hypothesis that halos have high central densities: not only does friction cause bars to become too slow; they also become unrealistically large and strong.

7.3. Disk Velocity Dispersion

A cooler disk leads to a bar of higher pattern speed that may also have to be shorter to fit inside its corotation circle.⁸ The trend in Ω_p is expected from bar instability theory (Kalnajs 1977) and agrees with previous findings (Athanassoula & Sellwood 1986) and with the difference between the models A_1 and A_2 reported by VK03.

Figure 10 presents an additional illustration of the effect of lower velocity dispersion at the time the bar formed; $Q \simeq 1.5$ in

⁸ A weak bar may have a lower Ω_p and end well inside its corotation circle.

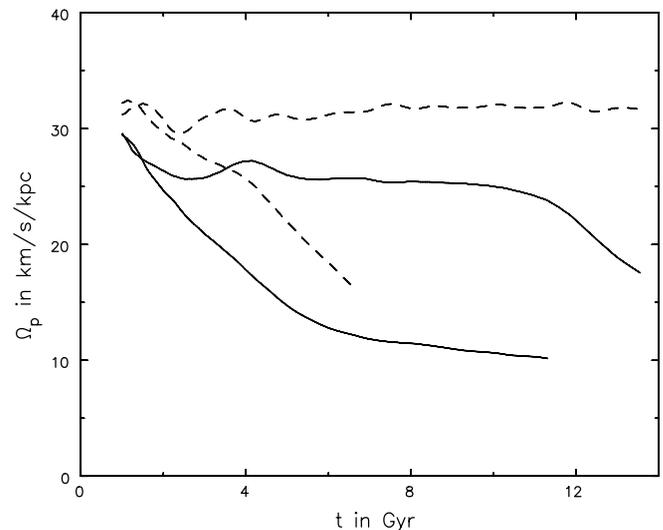


FIG. 10.—Time evolution of runs F and S, reproduced from Fig. 1 (*solid lines*). The dashed curves show results when the gravitational forces are restricted sectoral harmonics $0 \leq m \leq 4$.

the disks of the models with dashed lines, whereas $Q \simeq 1.8$ for the models with solid lines. The solid lines are from our runs F and S, reproduced from Figure 1, in which sectoral harmonics $0 \leq m \leq 8$ contributed to the forces computed from the disk particles. The dashed lines show results from models in which all sectoral harmonics $m > 4$ were eliminated from the force determination, i.e., a calculation with more heavily smoothed forces.

The disk is cooler in the models with the more restricted forces because the early stages of spiral evolution are weaker; this is not to say that strong spiral patterns with more than four arms develop when more harmonics are included, but simply that the patterns that do develop have steeper density gradients than simple sinusoidal density profiles, leading to stronger scattering by the density fluctuations.

The pattern speed clearly stayed fast in one of these two models for as long as we ran it, but a minor change to the numerical parameters (we increased the number of radial points of the spherical grid) again caused the bar to slow down. Note that Ω_p slowed at about the same rate as for the runs with a larger number of sectoral harmonics; the bar with the higher Ω_p is somewhat shorter at first, but grows to about the same size as it slows, causing the torques to become more nearly equal.

7.4. Collisionality

Figure 11 compares the angular momentum transferred from the disk to the halo in our run F (Fig. 11; *solid curve*) with that reported by VK03 for the same model; the dot-dash line is reproduced from their Figure 10. The disk in the simulation with their ART code loses angular momentum more rapidly than in ours. In both simulations, the torque is strong when Ω_p is changing and is weaker while Ω_p is approximately constant. However, VK03 find that the disk in their simulation starts to lose angular momentum right from the outset, before the bar has formed, and the gradient is again steeper than in our case while the bar rotates steadily.

The more rapid angular momentum transfer in their model during the period of steady bar rotation appears to be unrelated to the usual bar friction. As already noted in § 2, the properties of the initial bar in their simulation are similar to those in all the reruns reported here. We cannot compare our bar strengths with theirs, since we employ a different measure of bar strength, but

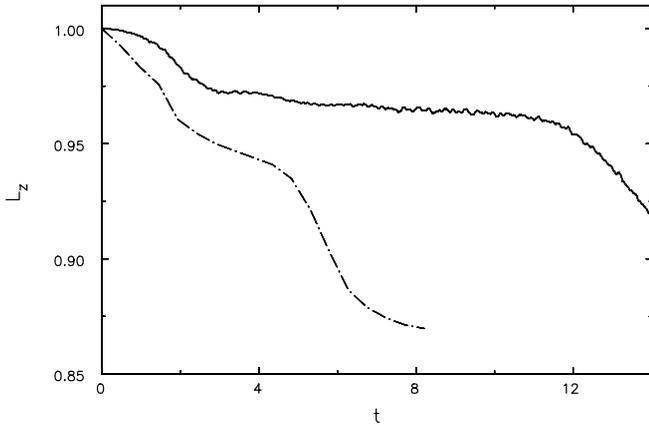


FIG. 11.—Disk angular momentum in our run F (*solid line*), relative to its initial value, as a function of time, showing that the torque is weak while the bar rotates steadily. The dot-dashed line, which is reproduced from Fig. 10 of VK03, shows the same quantity in their simulation.

initially similar bars in our simulations (Fig. 2) slow both more and less rapidly than in theirs, making it unlikely that the difference could be attributed to a stronger bar in their case.

It is possible that the stronger drag between the disk and halo in their simulation is caused by two-body encounters in their code. As a halo particle pursues its orbit, the density of nearby particles peaks up briefly as it passes through the disk. In an adaptive code that refines any grid cell that contains more than just a few particles, the likelihood of a strong deflection from a close encounter with another particle is greatly enhanced within the disk. Since the disk particles generally have a larger L_z than those in the halo, deflections caused by close encounters between them will, on average, transfer angular momentum from the disk to the halo.

VK03 report a measurement of the collision rate in their code in a model without a disk (their § 3.4), which is an ensemble average of all the encounters in high- and low-density regions. Their measurement is dominated by the cumulative effect of distant encounters.

Here we appeal exclusively to close encounters that produce the occasional strong deflection. It is easy to show that, to first order, deflections are just as likely to add angular momentum to disk particles as they are to remove it, and a net torque from collisions arises only in second order. Since the mean angular momentum exchanged increases as the square of the angle of deflection, the net collisional torque is dominated by rare, large-angle deflections.

We suspect that the nonzero disk-halo torque reported by VK03 from the outset of their experiments results from two-body encounters as halo particles pass through the disk. Note that the density of particles in the disk midplane is more than twice that of the halo alone over a wide swath of the disk, further increasing the likelihood of a strong deflection in this critical part of the orbit of a halo particle. Note also that collisional encounters cause angular momentum to be lost even when the disk is axisymmetric and, since they remove angular momentum from all over the disk and not just from the bar, the bar can rotate steadily while angular momentum continues to be lost to the halo. Of course, encounters must continue to scatter halo particles after the disk has developed a strong bar, adding to the torque that arises from dynamical friction.

Such encounters must also scatter the disk particles, and we therefore expect the disk to thicken more in their experiments than in ours. Their paper does not provide any information about disk thickness with which we might compare, however.

Collisionality could be an additional reason that the metastable state is short-lived in their simulation. Collisional encounters, whether between halo-halo particles, or halo-disk particles, would scatter particles and erode unusual density gradients near resonances, allowing friction to resume after less evolution than in a more nearly collisionless code.

8. CONCLUSIONS

The bars in two simulations reported by VK03 and those in two of our own reported here rotate in dense halos for long periods without the expected friction. We have shown that this anomalous behavior arises because the halos lack the decreasing density of halo particles with angular momentum about the principal resonances usually responsible for friction. An inflexion in the angular momentum density of particles is set up during a period of normal friction as the pattern speed decreases in the usual manner. A subsequent increase in the bar pattern speed can move the resonance into the region where the local gradient is anomalous, with the result that friction is suppressed for a long period. Because it relies on local minima in the distribution of halo particles, we describe almost frictionless bar rotation as a metastable state. It does not last indefinitely, although it can persist for cosmologically interesting time periods in isolated galaxy simulations.

We argue that the long period of steady bar rotation reported by VK03 in the evolution of their models A₁ and B occurred for this reason. We occasionally find that bars enter a metastable state in our own simulations, but we have shown that changes to the numerical procedure, or to numerical parameters, can cause or prevent the state from appearing. It arises in our models from a rare interaction between the bar and a stochastic spiral in the outer disk that gives the bar some angular momentum. Such an event could also have happened in the simulations by VK03, but the metastable state is more likely a result of their numerical method. We show that the pattern speed is driven upward by increases in the central attraction as the grid is refined in the adaptive scheme used by VK03, making the metastable state accessible for numerical reasons. It should be noted that VK03 report in their § 7 that bars slow normally in their code when adaptive refinement is turned off. Their model C, in which the level of refinement was also limited, was not run for long enough to show the full extent of bar slowdown.

We have shown that when this artifact is avoided, the bar in their model A₁ experiences strong friction and the corotation radius quickly moves out to an unacceptable distance, relative to the bar length. Thus, their model is, in fact, consistent with the conclusions of DS00 that strong bars cannot stay fast if the halo density is high.

We emphasize that bars can rotate at nearly constant pattern speed for reasons other than metastability. Friction will be mild and bars will not slow much if either the halo density is low or the bar is weak. The metastable state describes only strong bars that are able to rotate with little friction in dense halos because the distribution of resonant particles that would normally give rise to strong friction has acquired an anomalous gradient in angular momentum.

If strong bars in real galaxies could rotate rapidly in dense halos for long periods, the constraint on halo density proposed by us (DS98; DS00) would be severely weakened. In simulations without dissipation such behavior is a numerical artifact, but we have shown that mimicking a realistic amount of gas inflow in the bar can indeed cause the bar to speed up sufficiently for friction to stop, suggesting that the metastable state could arise quite naturally. However, we have also shown that the metastable

state is highly fragile and does not survive even minor internal or external perturbations. We find that a flyby satellite having as little as 1% of the main galaxy mass provides a sufficient disturbance to the system that normal friction resumes. It is therefore unlikely that many galaxies could survive in the metastable state for long periods. It is doubly unlikely that SB0 galaxies, the majority of galaxies for which \mathcal{R} is measured, are in this state, since they lack gas that can be driven into the center to raise the bar pattern speed.

We thank Anatoly Klypin and Octavio Valenzuela for stimulating discussions and for comments on a draft of this paper. They openly shared their results ahead of publication and provided the file of initial coordinates that we have used in these tests. We also thank an anonymous referee for a helpful report. This work was supported by grants to J. A. S. from NASA (NAG 5-10110) and from NSF (AST 00-98282). V. P. D. is supported by a Brooks Prize Fellowship in Theoretical Physics at the University of Washington.

APPENDIX

BAR MEASUREMENTS

We need to estimate the length of the bar from the simulations. As emphasized in previous studies, this is a nontrivial task because the bar has no sharp edges, and the bisymmetric density variations in the simulation often include spirals and, less often, mildly distorted rings. Our procedure is motivated in part by that followed by observers (e.g., Aguerri et al. 2003), who have a smooth light distribution instead of the noisy set of particles in our problem, but who have only a single snapshot of an inclined disk.

Since we use a polar grid for the disk component, it is a simple matter to save the amplitude and phase of the $m = 2$ component of the projected surface density on each radial ring at frequent intervals. We find the pattern speed of the bar, Ω_p , which we require to be the same at all radii, from a fit to these coefficients in the inner grid over a short time interval around the time at which we desire a measurement. We determine the corotation radius, R_c , from the azimuthally averaged rotation curve at the midpoint in time of the selected data. We have not tried to determine the Lagrange point, as did DS00, since the difference from the averaged corotation radius is generally small.

Each fit for Ω_p also yields the radial profile of the relative amplitude and phase of the $m = 2$ density variations, which are equivalent to aligning the data from each moment along a common axis dictated by the fitted Ω_p . Averaging in this way both diminishes shot noise from the particles and weakens the significance of any other bisymmetric feature, such as a spiral pattern, that may rotate at a different pattern speed. We use the ratio of the total amplitude of the $m = 2$, relative to the $m = 0$, terms from this time-averaged density as a measure of the bar amplitude.

The radius at which the relative amplitude of the $m = 2$ density is half that of its peak value is generally quite stable over time, but is clearly an underestimate of the bar length; unfortunately, we found that the radius where the amplitude is only 10% of the peak fluctuates quite wildly because of imperfectly eliminated spiral arms and the like. We therefore persisted with this low estimate (1), the radius of the half-peak amplitude, and sought a second high estimate (2), which is the radius at which the phase shifts by 20° from its mass-weighted average value near the peak. As this latter is generally quite clearly an overestimate of the bar length, we define the bar length, a_b , to be the average of estimates 1 and 2.

The dimensionless ratio $\mathcal{R} = R_c/a_b$, which can be compared with observed values (§ 1), is generally greater than unity. Unfortunately, estimate 2 of a_b suffers from large upward fluctuations at times, which are only half-eliminated by averaging and lead to corresponding downward fluctuations in \mathcal{R} , which are readily recognized because we monitor the value continuously—an example is visible in Figure 2. Our quoted values of \mathcal{R} generally ignore large, short-lived downward fluctuations. We have found that bars in cool, massive disks generally form with $\mathcal{R} \simeq 1$, as found in previous work, and that \mathcal{R} fluctuates around unity in an extreme maximum-disk model in which the bar was hardly braked at all, giving us confidence in our estimation method.

REFERENCES

- Aguerri, J. A. L., Debattista, V. P., & Corsini, E. M. 2003, *MNRAS*, 338, 465
 Athanassoula, E. 1992, *MNRAS*, 259, 345
 ———. 2002, *ApJ*, 569, L83
 ———. 2003, *MNRAS*, 341, 1179
 Athanassoula, E., & Sellwood, J. A. 1986, *MNRAS*, 221, 213
 Binney, J., & Tremaine, S. 1987, *Galactic Dynamics* (Princeton: Princeton Univ. Press)
 Buta, R., & Combes, F. 1996, *Fundam. Cosmic Phys.*, 17, 95
 Debattista, V. P., & Sellwood, J. A. 1998, *ApJ*, 493, L5 (DS98)
 ———. 2000, *ApJ*, 543, 704 (DS00)
 Debattista, V. P., & Williams, T. B. 2004, *ApJ*, 605, 714
 Hernquist, L. 1990, *ApJ*, 356, 359
 Hernquist, L., & Weinberg, M. D. 1992, *ApJ*, 400, 80
 Holley-Bockelmann, K., & Weinberg, M. 2005, *BAAS*, 36, 0512
 Kalnajs, A. J. 1977, *ApJ*, 212, 637
 Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, *ApJ*, 522, 82
 Kravtsov, A. V., Klypin, A., & Khokhlov, A. M. 1997, *ApJS*, 111, 73
 Lin, D. N. C., & Tremaine, S. 1983, *ApJ*, 264, 364
 Little, B., & Carlberg, R. G. 1991, *MNRAS*, 250, 161
 Moore, B., Governato, F., Quinn, T., Stadel, J., & Lake, G. 1998, *ApJ*, 499, L5
 Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, *ApJ*, 490, 493
 O’Neill, J. K., & Dubinski, J. 2003, *MNRAS*, 346, 251
 Rautiainen, P., Salo, H., & Laurikainen, E. 2005, *ApJ*, 631, L129
 Sakamoto, K., Okamura, S. K., Ishizuki, S., & Scoville, N. Z. 1999, *ApJ*, 525, 691
 Sellwood, J. A. 1981, *A&A*, 99, 362
 ———. 2003, *ApJ*, 587, 638
 ———. 2006, *ApJ*, 637, 567 (Paper I)
 Sellwood, J. A., & Merritt, D. 1994, *ApJ*, 425, 530
 Toomre, A. 1981, in *The Structure and Evolution of Normal Galaxies*, ed. S. M. Fall & D. Lynden-Bell (Cambridge: Cambridge University Press), 111
 Tremaine, S., & Weinberg, M. D. 1984, *MNRAS*, 209, 729
 Valenzuela, O., & Klypin, A. 2003, *MNRAS*, 345, 406 (VK03)
 Weinberg, M. D. 1985, *MNRAS*, 213, 451
 ———. 2004, preprint (astro-ph/0404169)
 Weinberg, M. D., & Katz, N. 2002, *ApJ*, 580, 627
 ———. 2005, *MNRAS*, submitted (astro-ph/0508166)