



Rare Event Sampling Improves Mercury Instability Statistics

Dorian S. Abbot¹, Robert J. Webber², Sam Hadden³, Darryl Seligman¹, and Jonathan Weare²¹Department of the Geophysical Sciences, The University of Chicago, Chicago, IL 60637 USA; abbot@uchicago.edu²Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 USA³Harvard-Smithsonian Center for Astrophysics, The Institute for Theory and Computation, Cambridge, MA 02138, USA

Received 2021 June 8; revised 2021 September 30; accepted 2021 October 11; published 2021 December 24

Abstract

Due to the chaotic nature of planetary dynamics, there is a non-zero probability that Mercury’s orbit will become unstable in the future. Previous efforts have estimated the probability of this happening between 3 and 5 billion years in the future using a large number of direct numerical simulations with an N -body code, but were not able to obtain accurate estimates before 3 billion years in the future because Mercury instability events are too rare. In this paper we use a new rare-event sampling technique, Quantile Diffusion Monte Carlo (QDMC), to estimate that the probability of a Mercury instability event in the next 2 billion years is approximately 10^{-4} in the REBOUND N -body code. We show that QDMC provides unbiased probability estimates at a computational cost of up to 100 times less than direct numerical simulation. QDMC is easy to implement and could be applied to many problems in planetary dynamics in which it is necessary to estimate the probability of a rare event.

Unified Astronomy Thesaurus concepts: [Planetary dynamics \(2173\)](#)

1. Introduction

Laskar (1994) used secular simulations, with the equations of motion averaged over planetary orbits, to reach the surprising conclusion that it is possible for Mercury’s orbit to become unstable. The mechanism underlying Mercury’s potential instability involves resonances among the secular system’s modes of oscillation that can transfer angular momentum among the planets and dramatically increase Mercury’s eccentricity (Lithwick & Wu 2011; Boué et al. 2012; Lithwick & Wu 2014; Batygin et al. 2015). As a result, Mercury can pass very near Venus, disrupting its orbit, and leading to a collision with the Sun or another planet. Laskar (2008) performed a large number of secular simulations with slightly different initial conditions to estimate that the probability of this happening is $\mathcal{O}(10^{-2})$ in the next 5 billion years (Gyr). Laskar & Gastineau (2009) and Zeebe (2015a) each performed $\mathcal{O}(10^3)$ N -body simulations that roughly confirmed the secular results. Specifically, we can use the results from Laskar & Gastineau (2009) to calculate the probability of a Mercury instability event in the next 5 Gyr as $8.0 \pm 1.8 \times 10^{-3}$ and those from Zeebe (2015a) to calculate it as $6.3 \pm 2.0 \times 10^{-3}$, assuming 1σ errors calculated using Equation (1).

Since a Mercury instability event would represent a major milestone in Solar System evolution, it is interesting to consider the probability of such events in comparison to other milestones. Laskar (1994) chose a 5 Gyr time frame because that is the remaining main-sequence lifetime of the Sun. Another key milestone will occur in roughly 2 Gyr when Earth loses its habitability either through water loss or a runaway greenhouse (Wolf & Toon 2015), which motivates trying to evaluate the probability of Mercury instability events at earlier times. Although the simulations of Laskar & Gastineau (2009) and Zeebe (2015a) allow us to calculate the probability of

Mercury instability events in the next 5 Gyr, these events are too rare in the next 3 Gyr for their probability to be estimated accurately from their simulation suites. Specifically, in the first 3 Gyr there are only two occurrences out of 2501 simulations performed by Laskar & Gastineau (2009) and only one occurrence out of 1600 simulations performed by Zeebe (2015a). It would be possible to improve event statistics by greatly increasing the number of simulations performed, but the associated computational cost makes this an undesirable solution.

In this paper we will use a rare-event sampling technique to obtain accurate estimates of the probability that Mercury’s orbit becomes unstable in the next 2 to 3 Gyr without increasing the computational cost. This method is called Quantile Diffusion Monte Carlo (QDMC; Webber et al. 2019). Using QDMC we pause the simulations regularly and check a “reaction coordinate” that indicates progress toward the desired event. We then rank the simulations in terms of this reaction coordinate. Next we preferentially restart multiple copies of simulations whose rank progressed more toward the desired event with slightly different initial conditions (splitting) and preferentially stop simulations whose rank progressed less toward the desired event (killing). We both split and kill probabilistically, and throughout this process we carefully track the splitting and killing history so that we can recover the statistics of the underlying model. As we do this, we maintain a constant total number of simulations and therefore constant computational cost. The net effect of QDMC is that we devote a vastly disproportionate amount of computational power to simulating the rare event of interest, and as a result greatly improve our estimate of its probability and the uncertainty in this probability. Using QDMC, we are able to obtain unbiased estimates of the probability of a Mercury instability event between 2 and 3 Gyr in the future at a computational cost that is up to 100 times less than direct numerical simulation.

Our work builds on two previous efforts in this area that have also used splitting and killing (Laskar 1994; Batygin & Laughlin 2008). Critically, both of those efforts merely split the simulation that had progressed the most and killed the others,



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

which prevented them from reconstructing event statistics in the underlying model. Moreover, the underlying model of Laskar (1994) was secular and the model of Batygin & Laughlin (2008) did not include general relativity, which has a dramatic effect on Mercury instability event statistics (Laskar 2008; Laskar & Gastineau 2009). In this paper we will use the Rebound N -body code (Rein & Liu 2012), including expansion terms that represent general relativity (Tamayo et al. 2020). Finally, both Laskar (1994) and Batygin & Laughlin (2008) used Mercury’s eccentricity to measure progress toward Mercury instability events (as a reaction coordinate). We will show that Mercury’s instantaneous or time-averaged eccentricity is a poor predictor of Mercury instability events ~ 0.5 Gyr in the future. Instead, we use the range (range = max – min) of the Mercury–Venus Minimum Orbit Intersection Distance (MOID) 0.4–0.8 Gyr before a potential instability event as our reaction coordinate. The Mercury–Venus MOID range can be calculated quickly and is correlated with the Mercury–Venus orbital energy. We find that large variations in the MOID are good predictors of future instability in Mercury’s orbit, possibly because they are a sign of energy being pumped into and out of Mercury’s orbit by other planets.

This paper is organized as follows. We describe our methods in Section 2, including the REBOUND N -body code that we will use (Section 2.1), our direct numerical simulations (Section 2.2), our implementation of QDMC (Section 2.3), and our choice of reaction coordinate (Section 2.4). We give our results in Section 3, including showing that our direct numerical simulation results are consistent with previous work (Section 3.1) and demonstrating the improvement in Mercury instability event statistics we can obtain with QDMC (Section 3.2). We discuss the limitations and implications of our work in Section 4 and conclude in Section 5.

2. Methods

Our implementation of QDMC for REBOUND as well as the data and code to produce all of the figures in this paper are permanently archived in the Knowledge@UChicago repository at the URL <https://knowledge.uchicago.edu/record/3323?&ln=en>.

2.1. REBOUND

We perform all simulations in this paper using the REBOUND N -body code’s (Rein & Liu 2012) WHCKL integration scheme (Rein et al. 2019), which applies symplectic correctors (Wisdom 2006) and a modified kick step (Wisdom et al. 1996) to a classic Wisdom–Holman scheme (Wisdom & Holman 1991) in order to achieve an error scaling of $\mathcal{O}(\epsilon dt^{18} + \epsilon^2 dt^4 + \epsilon^3 dt^3)$ where $\epsilon \sim 10^{-3}$ in the case of the solar system. We adopt a fixed time step of $dt = \sqrt{65} \approx 8.062$ days. This is small enough to ensure that Mercury’s orbit is properly resolved and irrational in order to avoid step-size resonances. Our integrator configuration matches the configuration used for the integrations described in Brown & Rein (2020). Our integrations contain all eight Solar System planets and include an approximation of general relativity (Tamayo et al. 2020). We start the simulations from current Solar System conditions with a Gaussian random perturbation to Mercury’s x-coordinate position with a length scale of 1 cm. We stop the simulations if Mercury and Venus come within 0.01 au of each

other (the length scale of Venus’s Hill sphere), which we use to define a Mercury instability event. This is a slightly more inclusive definition of a Mercury instability event than those used by Laskar & Gastineau (2009) and Zeebe (2015a), since they traced Mercury until a collision occurred.

As described by Zeebe (2015b), there are many challenging numerical issues associated with the performance of N -body integrators. These issues become especially apparent when planets’ eccentricities become high or when they approach each other. For example, the symplectic method can become unstable and may introduce artificial chaos unless the time step is small enough to always resolve perihelion (Wisdom 2015; Zeebe 2015b). The REBOUND configuration we are using should be accurate until shortly before a Mercury instability event occurs, which is consistent with the fact that our direct numerical simulation estimates of the probability of Mercury instability events are similar to those of Laskar & Gastineau (2009) and Zeebe (2015a) (Section 3.1). In any case, what is most important for this paper is that QDMC can produce unbiased results at reduced computational cost compared to direct numerical simulation with the same underlying model.

2.2. Direct Numerical Simulation

We perform 1008 direct numerical simulations of the Solar System for 5 Gyr. In order to make the underlying model of the direct numerical simulations exactly the same as the QDMC simulations described in Section 2.3, we apply a Gaussian random perturbation to Mercury’s position with a length scale of 1 cm every 0.2 Gyr in our direct numerical simulations. Given the chaotic nature of Solar System dynamics (Laskar 1989), we do not expect these perturbations to change the statistics of Mercury instability events. They may, however, affect which particular simulations experience a Mercury instability event.

Our different Mercury simulations evolve independently, so the number of observed instability events is a binomial random variable. Let p_e denote the true probability of instability, N the number of simulations, and N_e the number of observed instability events. We estimate p_e using the maximum likelihood formula $\hat{p}_e = N_e/N$, where hat denotes an estimated quantity, and estimate 1σ error bars using

$$\hat{\sigma}_p = \sqrt{\frac{\hat{p}_e(1 - \hat{p}_e)}{N}}. \quad (1)$$

We will use the estimates \hat{p}_e and $\hat{\sigma}_p$ for our direct numerical simulation results, as well as those of Laskar & Gastineau (2009) and Zeebe (2015a).

2.3. QDMC

Background. The basic principle underlying QDMC is that an ensemble of simultaneous model simulations can be pushed toward a rare event of interest by occasional duplication of some ensemble members and removal of others (resampling). Incarnations of this “splitting” idea appear in print at least as early as the 1950s (Kahn & Harris 1951; Rosenbluth & Rosenbluth 1955). Since then, many splitting schemes have been used in rare-event simulation problems in physics (Grassberger 1997), chemistry (Huber & Kim 1996; Allen et al. 2006; Warmflash et al. 2007; Guttenberg et al. 2012), computer science (Villen-Altamirano & Villen-Altamirano 1991;

Haraszti & Townsend 1999), statistics (Moral & Garnier 2005; Cérou & Guyader 2007), and geophysics (Ragone et al. 2018; Webber et al. 2019; Ragone & Bouchet 2021).

The first and most important step in the application of a rare-event simulation technique is the selection of a reaction coordinate. The reaction coordinate quantifies the progress of an ensemble member toward the rare event, and the change in its value since the last resampling determines the likelihood that the ensemble member will be duplicated or removed. In most situations the optimal reaction coordinate for computing the probability of a rare event is the probability of the event as a function of the current state of the system (e.g., Dupuis & Wang 2004; Dean & Dupuis 2009, 2011; Vanden-Eijnden & Weare 2012; Webber et al. 2020). However, in many cases a very rough descriptor of this probability is sufficient. In applications of real scientific interest, determination of an effective reaction coordinate often requires a non-trivial computational interrogation of the system (Antoszewski et al. 2020).

Quantile DMC (QDMC, Webber et al. 2019) is a recently proposed splitting method that builds on previous approaches by improving the robustness properties. As opposed to other splitting and killing schemes (e.g., Giardina et al. 2006; Ragone et al. 2018), QDMC is based on the sorting of trajectories along a reaction coordinate. First, QDMC sorts simulations from the lowest reaction coordinate value to the highest reaction coordinate value. Next, QDMC randomly splits simulations that have moved from lower to higher relative positions in the reaction coordinate and randomly kills simulations that have moved from higher to lower relative positions in the reaction coordinate. Because QDMC uses relative positions along a reaction coordinate instead of absolute positions, it is more robust to the choice of reaction coordinate (Webber et al. 2019).

Algorithmic details. QDMC starts with an initialization step and then iterates over a quantile transformation step, a splitting and killing step, and a forward evolution step, as detailed below.

1. *Initialization.* Generate independent simulations $X_{t_0}^1, \dots, X_{t_0}^N$ and assign equal weights to each simulation, i.e., $w_{t_0}^i = 1/N$ for each $i = 1, \dots, N$.
2. *Iteration.* Apply steps (a)–(c) at uniformly spaced times $t = t_0, t_0 + \Delta, t_0 + 2\Delta, \dots$
 - (a) *Quantile transformation.* Evaluate the reaction coordinate $\theta(X_t^i)$ for each simulation X_t^1, X_t^2, \dots , and find a permutation $\alpha(1), \alpha(2), \dots$ that reorders the reaction coordinates from lowest to highest, i.e.,

$$\theta(X_t^{\alpha(1)}) \leq \theta(X_t^{\alpha(2)}) \leq \dots \quad (2)$$

Next, for each simulation X_t^1, X_t^2, \dots , define a rescaled reaction coordinate

$$\hat{\theta}_t(X_t^i) = \Phi^{-1} \left(\sum_{j: \theta(X_t^j) < \theta(X_t^i)} w_t^j + \frac{w_t^i}{2} \right), \quad (3)$$

where Φ^{-1} is the inverse of the Gaussian cumulative distribution function $\Phi(x) = \int_{-\infty}^x (2\pi)^{-1/2} \exp(-z^2/2) dz$. This is done to ensure that the rescaled reaction coordinates are approximately Gaussian distributed.

- (b) *Splitting and killing.* Divide the simulations into stable and unstable trajectories. Set aside the unstable trajectories for statistical analysis later: these

trajectories remain part of the statistical sample but their forward evolution is stopped. Split each stable simulation X_t^i into N_t^i identical replicas, where the numbers N_t^i are randomly chosen to satisfy

$$\mathbb{E}N_t^i = \frac{Nw_t^i \exp(k_t \hat{\theta}_t(X_t^i))}{\sum_{j: X_t^j \text{ is stable}} w_t^j \exp(k_t \hat{\theta}_t(X_t^j))} \quad (4)$$

and $\sum_{i: X_t^i \text{ is stable}} N_t^i = N,$

for an intensity parameter $k_t > 0$. Assign the children of X_t^i weights summing to w_t^i in expectation, following the procedure in Appendix.

- (c) *Forward evolution.* Advance the stable trajectories forward until the next resampling time, using independent random seeds.
3. *Statistical estimation.* At any time, use the weighted measure $\sum_i w_t^i \delta_{X_t^i}$ to produce unbiased estimates with respect to the simulation model, e.g.,

$$\text{Prob}\{\text{instability at time } t\} \approx \sum_{i: X_t^i \text{ is unstable}} w_t^i. \quad (5)$$

The sum of the weights $\sum_i w_t^i$ is exactly one, and these weights provided the estimated probabilities of different events occurring.

In QDMC, the dynamics are determined by a rescaled reaction coordinate $\hat{\theta}$ and by an intensity function $k_t > 0$ (in standard deviation units), which controls the strength of the splitting and killing. Using direct numerical simulations, the distribution of $\hat{\theta}_t(X_t^1), \hat{\theta}_t(X_t^2), \dots$ would be nearly $\mathcal{N}(0, 1)$. However, QDMC drives the distribution of rescaled reaction coordinates $\hat{\theta}_t(X_t^1), \hat{\theta}_t(X_t^2), \dots$ to nearly $\mathcal{N}(k_t, 1)$. In this sense, QDMC pushes the distribution of reaction coordinate values k_t standard deviation units higher, as compared to the underlying simulation model.

Parameter choices. In typical QDMC applications, we gradually increase the intensity function k_t leading up to a target time $t = t_{\text{target}}$. To obtain an explicit formula for k_t , we model the rescaled reaction coordinate $\hat{\theta}_t$ as a Gaussian linear process

$$\begin{cases} d\hat{\theta}_t = -(\hat{\theta}_t/\tau)dt + \sqrt{2/\tau}dW, & t > 0 \\ X_0 = 0, \end{cases} \quad (6)$$

with a decorrelation timescale τ , and we optimize k_t to yield the minimal variance estimate for $\mathbb{E} \exp(k_{\text{max}} X_{t_{\text{target}}})$ (see, e.g., Cérou & Guyader 2007; Webber et al. 2020). The resulting intensity function is

$$k_t = k_{\text{max}} \exp \left(\frac{t - t_{\text{target}}}{\tau} \sqrt{1 - \exp \left(\frac{-2t}{\tau} \right)} \right). \quad (7)$$

The tuning parameters for QDMC are thus the reaction coordinate k_t , the decorrelation timescale τ , the intensity threshold k_{max} , and the resampling times leading up to t_{target} .

QDMC produces unbiased, convergent estimates regardless of the specific parameters, yet the choice of parameters does affect QDMC's efficiency at converging to the correct statistics. The main parameter determining QDMC's efficiency is the reaction coordinate θ_t , which is discussed at length in Section 2.4. QDMC is comparatively less sensitive to the other

parameters, and we follow the recommendations given in Webber et al. (2019) as described below:

1. The decorrelation timescale τ should be tuned based on the timescale for large random changes in the rare-event probability. Since Mercury’s eccentricity can increase from .3 to .8 over a timescale of .5 Gyr (Laskar & Gastineau 2009; Zeebe 2015a) and since small Gaussian perturbations can potentially push Mercury toward or away from such spiraling instability, we set $\tau = .5$ Gyr. However, as explained in Section IIE of Webber et al. (2019), QDMC is not very sensitive to the decorrelation timescale, and over-estimating or under-estimating the timescale by a factor of two causes less than a 20% increase in QDMC’s error.
2. The intensity threshold k_{\max} should be tuned based on the approximate magnitude of the rare-event probability. Because the simulation results of Laskar & Gastineau (2009) and Zeebe (2015a) suggest the probability of Mercury becoming unstable 2–3 Gyr into the future is between 10^{-4} and 10^{-3} , we set the intensity threshold to be $k_{\max} = 3$, which is appropriate for estimating rare events with probabilities ranging from 3×10^{-5} – 2×10^{-2} (i.e., rare excursions of 2–4 standard deviations away from the mean).
3. There should be a modest number of resampling times (≤ 10) during the interval $[t_{\text{target}} - 2\tau, t_{\text{target}}]$ so that QDMC achieves its maximum efficiency at the target time t_{target} . Here, we use the specific resampling times

$$\begin{aligned} t_{\text{target}} - 1 \text{ Gyr}, & \quad t_{\text{target}} - .8 \text{ Gyr}, & \quad t_{\text{target}} - .6 \text{ Gyr}, \\ t_{\text{target}} - .4 \text{ Gyr}, & \quad t_{\text{target}} - .2 \text{ Gyr}, \end{aligned} \quad (8)$$

and we run three separate QDMC trials using $t_{\text{target}} = 2.4$ Gyr, $t_{\text{target}} = 2.8$ Gyr, and $t_{\text{target}} = 3.2$ Gyr. This schedule is slightly different from the recommendation given in (Webber et al. 2020), which suggests increasing the frequency of resampling times as $t \rightarrow t_{\text{target}}$. However, here we are limited in how short we can make the intervals between resampling times because it takes approximately 0.1–0.15 Gyr for orbits with small initial perturbations to spread apart in eccentricity values (Figure 2). To ensure adequate statistical variation among our samples, we apply splitting and killing at regularly spaced 0.2 Gyr intervals.

Variance estimation. We provide two variance estimators for QDMC, one which overestimates the variance and another which underestimates the variance. For any statistical estimate of the form $\hat{f} = \sum_i w_i^j f(X_i^j)$, we can overestimate the variance using the following *pessimistic* variance estimator

$$\hat{\sigma}_{\text{pess}}^2 = \sum_i \left| \sum_{\text{anc}(X_i^j)=i} w_i^j f(X_i^j) \right|^2 - \hat{f}^2. \quad (9)$$

Here, $\text{anc}(X_i^j)$ denotes the index for the original ancestor of the simulation X_i^j traced back to time $t = 0$. We can underestimate the variance using the following *optimistic* variance estimator

$$\hat{\sigma}_{\text{opt}}^2 = \sum_i |w_i^i f(X_i^i)|^2 - \hat{f}^2. \quad (10)$$

To intuitively explain the difference between these two variance estimators, we remark that the optimistic estimator treats each individual simulation as a single data point as in

importance sampling (Liu 2008), while the pessimistic estimator treats each family of simulations as a single data point. However, the true sample size lies somewhere between these two extremes. Simulations with a lot of shared ancestry are often quite correlated with one another; however, simulations that diverged many generations ago are hardly correlated at all. While this explanation captures the essential intuition, we note that variance estimation for splitting schemes remains an open area of mathematical investigation. As of the present writing, the pessimistic variance estimator is known to overestimate the variance for large sample sizes (Webber et al. 2019), and the optimistic variance estimator is found empirically to underestimate the variance (Webber 2022). Further work is needed to close the gap between these two variance estimators.

Once we have obtained an estimate of the probability of a Mercury instability event and the associated error using QDMC, we can estimate the number of direct numerical simulations that would be necessary to produce a similarly accurate estimate using Equation (1) as follows

$$N \approx \frac{\hat{p}_e(1 - \hat{p}_e)}{\hat{\sigma}_p^2}. \quad (11)$$

Since all of our QDMC implementations involve 1008 simulations, we can estimate the speed-up provided by QDMC as $\frac{N}{1008}$, where N is given by Equation (11).

Improvements over past work. Compared to our previous application of QDMC to compute intensity statistics for tropical cyclones (Webber et al. 2019), we note two improvements. First, following recent research on the stability of splitting and killing schemes (Webber et al. 2020), we constrain the sum of the weights $\sum_i w_i^j$ to be exactly one, which promotes greater interpretability and ensures stability in the limit as $t \rightarrow \infty$. Second, while before we only used the pessimistic variance estimator $\hat{\sigma}_{\text{pess}}^2$ to compute error bars, here we also provide the optimistic variance estimator $\hat{\sigma}_{\text{opt}}^2$, and we compare the two estimators.

2.4. Choice of Reaction Coordinate

The ideal reaction coordinate would encompass all early warning signs that a Mercury instability event may occur. Identifying such an ideal reaction coordinate is extremely difficult; however, we can search for a substitute reaction coordinate that is easy to compute and indicates the likely occurrence of a Mercury instability event. In this work, we consider the median and range (range = max – min) of the Mercury–Venus MOID (measured in au), the Venus–Earth MOID, and Mercury’s eccentricity as potential reaction coordinates.

To assess the usefulness of potential reaction coordinates, we analyze our direct numerical simulation data set. This data set encompasses 996 trajectories that remained stable throughout the entire 5 Gyr integration and 12 trajectories in which a Mercury instability event occurred. We consider times 400–800 Myr before the simulation ended either at 5 Gyr or in a Mercury instability event. Then, using the time series for the Mercury–Venus MOID, the Venus–Earth MOID, and Mercury’s eccentricity, we compute the median and the range for statistical analysis.




Figure 1. The Mercury–Venus MOID is the most effective predictor of Mercury instability events 400–800 Myr in the future. Barplot shows correlation coefficients between Mercury stability/instability and six potential reaction coordinates, calculated 400–800 Myr before the simulation end. The potential reaction coordinates are the median and range of the Mercury–Venus MOID, the Venus–Earth MOID, and Mercury’s eccentricity.

As a first statistical test, we calculate correlation coefficients between the binary variable indicating stability/instability and the six predictor variables, as shown in Figure 1. This comparison suggests that the Mercury–Venus MOID range and Mercury’s eccentricity range are the strongest instability predictors, with correlation coefficients of $r = .26$ and $r = .24$, respectively. The medians of the Mercury–Venus MOID and Mercury eccentricity are comparatively poor predictors of instability ($r = -.10$ and $r = .10$), while the median and range of the Venus–Earth MOID have nearly no predictive potential ($r = .01$ and $r = .01$). Note that this implies that a simulation with large variation in its Mercury–Venus MOID is typically more likely to experience a Mercury instability event 400–800 Myr in the future than a simulation that has low Mercury–Venus MOID without much variation.

As a second statistical test, we run a Lasso logistic regression (Tibshirani 1996; Pedregosa et al. 2011) to predict future instability as a sparse linear combination of the six predictor variables. With a regularization parameter of $C = 0.01$, the resulting model is

$$\log\left(\frac{p}{1-p}\right) = -.224 + .569 X_{\text{MOID}} + .159 X_{\text{eccen}}, \quad (12)$$

where p is the probability of instability, while X_{MOID} and X_{eccen} represent the Mercury–Venus MOID range and Mercury’s eccentricity range, both normalized to have a mean of zero and a standard deviation of one. This model suggests that Mercury–Venus MOID range heavily influences the probability of instability, while Mercury’s eccentricity range makes a smaller secondary contribution.

As a result of this analysis, we use the Mercury–Venus MOID range calculated over 400 Myr of data as our reaction coordinate. Although we could potentially obtain better results by combining MOID data and eccentricity data into a single instability predictor, Equation (12) suggests that the role played by eccentricity in the optimal predictor is likely to be comparatively small. Using MOID data alone also leads to greater interpretability.

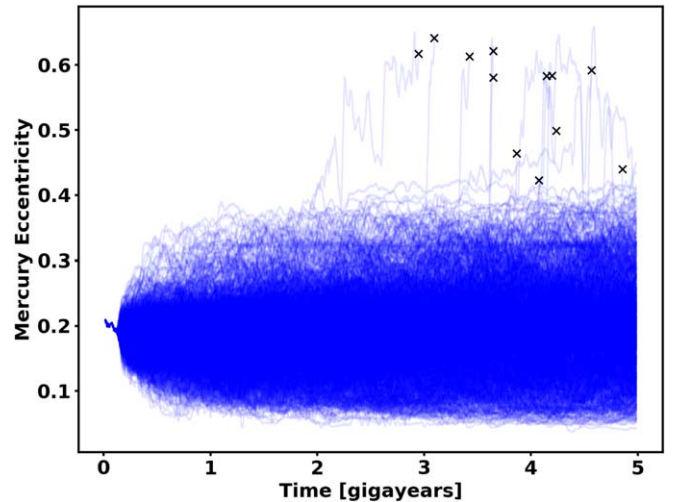


Figure 2. Mercury’s eccentricity from 1008 Direct Numerical Simulation (DNS) 5 Gyr integrations of the solar system. The eccentricity plotted is a 30 Myr running average. We add perturbations to Mercury’s position on the scale of 1 cm every 200 Myr in these integrations for similarity with our Quantile Diffusion Monte Carlo (QDMC) simulations, but this should not affect event statistics. The x symbols signify a close encounter between Mercury and Venus, at which point we stop our simulations. The x symbols occur in this plot when Mercury’s eccentricity is relatively small because of the running average and our output frequency.

We do not have a complete explanation for why the MOID range is an effective predictor of future Mercury instability events. Here we simply note that the Mercury–Venus MOID is strongly correlated with the Mercury–Venus orbit-averaged interaction energy. Given this feature of the MOID and the fact that Mercury’s eccentricity excitation is known to result from secular (i.e., orbit-averaged) dynamical interactions in which Venus plays a prominent role (e.g., Lithwick & Wu 2011), it is reasonable that the MOID range correlates with instability probability. The mechanism may be that large variations in the Mercury–Venus orbit-averaged interaction energy, captured by large variations in the Mercury–Venus MOID, are a sign of energy being pumped into and out of the system in a way that can lead to future instability.

3. Results

3.1. Direct Numerical Simulation

The paths of Mercury’s eccentricity for the next 5 Gyr in our 1008 direct numerical simulations are shown in Figure 2. We apply a 30 Myr running average to Mercury’s eccentricity in this plot so that individual simulations can be discerned. All of the simulations stay very near to each other for the first ~ 0.1 Gyr, at which point they start to diverge. For the most part, the spread in Mercury’s eccentricity increases slowly and roughly diffusively (Laskar 2008). Some simulations jump out of this slow spread in Mercury’s eccentricity, and eventually result in a close encounter between Mercury and Venus. Most of these Mercury instability events happen suddenly, with very little advance warning from Mercury’s eccentricity. Since a reaction coordinate must predict a rare event long in advance of the event occurring, this explains why Mercury’s eccentricity (rather than MOID range or eccentricity range) is not a good reaction coordinate for Mercury instability events.

The probability of Mercury instability events as a function of time in our direct numerical simulations is broadly similar to

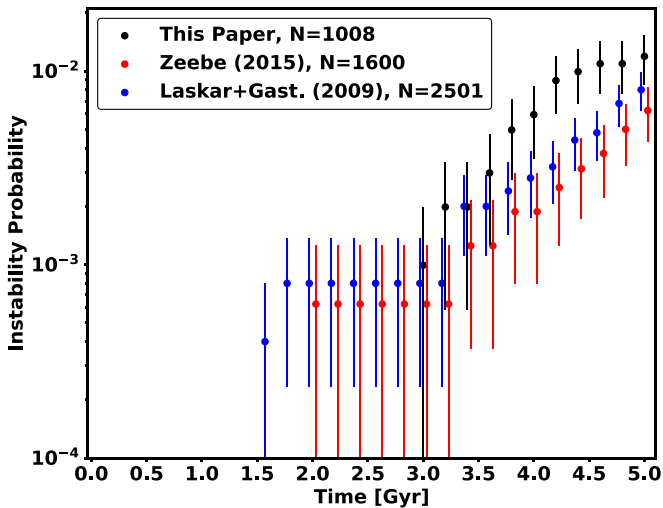


Figure 3. Our Direct Numerical Simulation (DNS) results are similar to previous studies, but slightly higher. An estimate of the probability that Mercury has a close encounter with Venus is plotted as a function of time for our DNS simulations (black), the DNS simulations of Zeebe (2015a) (red), and the DNS simulations of Laskar & Gastineau (2009) (blue). The error plotted is the 1σ error calculated using Equation (1).

that of Laskar & Gastineau (2009) and Zeebe (2015a) (Figure 3). For all three simulation suites the probability is lower than 10^{-3} before 3 Gyr in the future, then grows slowly to $\mathcal{O}(10^{-2})$ at 5 Gyr in the future. The probabilities we estimate are slightly higher than those of Laskar & Gastineau (2009) and Zeebe (2015a), especially between 3.5 and 4.5 Gyr in the future. This could be due to our more generous definition of Mercury instability events (Section 2.1), the fact that we do not decrease the time step at high eccentricity, and/or the chaotic nature of the system. This issue does not affect the primary goal of this project, which is to demonstrate that QDMC provides a computationally efficient unbiased estimate of Mercury instability statistics within a given model.

3.2. QDMC

QDMC produces many Mercury instability events at times before direct numerical simulation has produced any (Figure 4). In Figure 4 we can see the effects of splitting and killing, particularly as the simulation progresses and we force the model more strongly toward Mercury instability events. When viewing Figure 4, it is important to remember that we are not using Mercury’s eccentricity as our reaction coordinate. This is why extensive splitting can occur at relatively low values of Mercury’s eccentricity.

We plot our QDMC estimates of the Mercury instability event probability as a function of time in Figure 5. We can compare the QDMC results against direct numerical simulation estimates at 3.0 and 3.2 Gyr in the future, and the results are consistent within the standard error bars. Moreover, at times when we have QDMC implementations with different target times, the results are consistent within the error bars as well. At 2.4 Gyr all three of our QDMC implementations provide overlapping probability estimates. All of this is consistent with QDMC providing an unbiased estimate of the underlying Mercury instability probability of the model.

Using Equation (11) we can estimate the computational speed-up that QDMC provides relative to direct numerical simulation (Figure 6). The speed-up is higher around 2 Gyr in

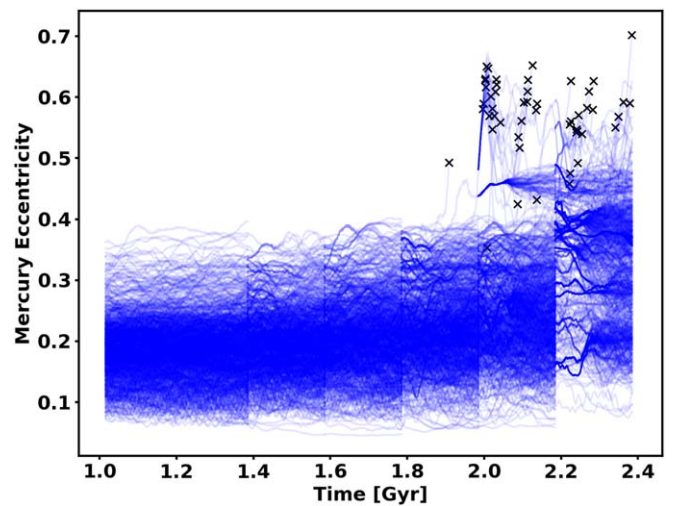


Figure 4. Quantile Diffusion Monte Carlo (QDMC) produces many more close encounters between Mercury and Venus than Direct Numerical Simulation (DNS). Here we show our results for QDMC where splitting occurs every 0.2 Gyr, starts at 1.4 Gyr, and ends at a target time of 2.4 Gyr. The vertical axis is Mercury’s eccentricity and the 30 Myr running average of the simulations is plotted. The x symbols signify a close encounter between Mercury and Venus.

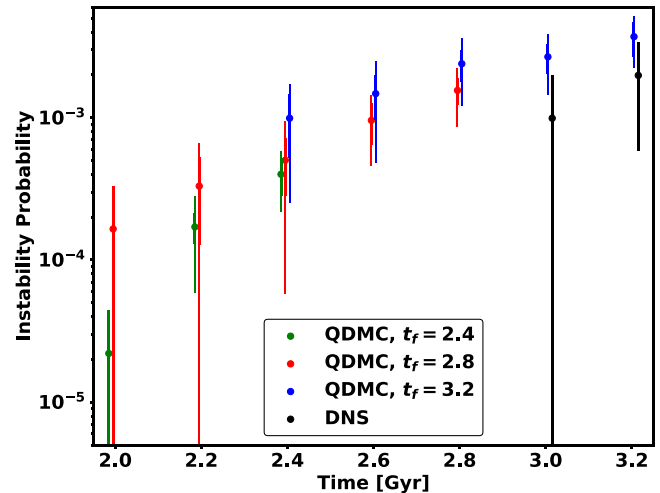


Figure 5. Quantile Diffusion Monte Carlo (QDMC) produces unbiased probability estimates. Here we are plotting the probability that Mercury has a close encounter with Venus as a function of time for our direct numerical simulations (DNS, black), our QDMC simulations with a target time of 2.4 Gyr (green), our QDMC simulations with a target time of 2.8 Gyr (red), and our QDMC simulations with a target time of 3.2 Gyr (blue). The 1σ error for each data point is also plotted. We calculated the error using Equation (1) for DNS. We provide an optimistic (thick) and pessimistic (thin) error estimate for QDMC, as described in Section 2.3.

the future than around 3 Gyr, which makes sense because QDMC should be more beneficial the rarer the event, and Mercury instability events are rarer at earlier times. Near 2 Gyr in the future, the speed-up reaches a maximum of about 100, assuming optimistic error bars. Even assuming pessimistic error bars, the speed-up is always greater than 10 for a QDMC target time of 2.4 Gyr. At QDMC target times of 2.8 Gyr and 3.2 Gyr the speed-up becomes progressively smaller, although it is still $\mathcal{O}(10)$ for optimistic error bars. Even assuming pessimistic error bars, QDMC always provides at least some speed-up relative to direct numerical simulation.

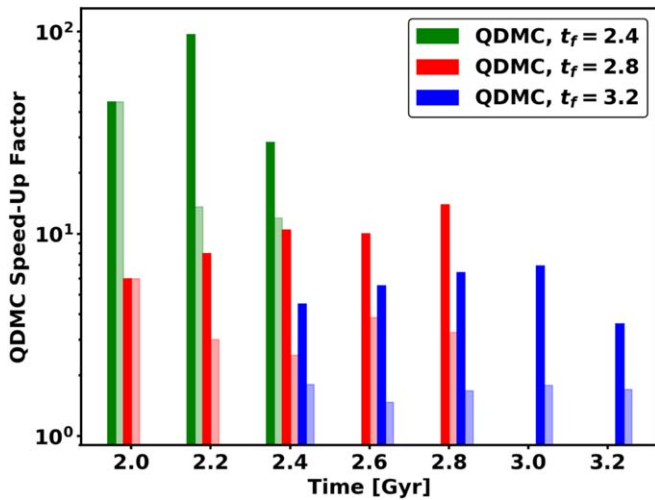


Figure 6. Quantile Diffusion Monte Carlo (QDMC) is up to 100 times faster than direct numerical simulation. This plot shows the ratio of the inferred number of direct numerical simulations needed for the precision of QDMC (Equation (11)) to the number used in our implementation of QDMC (1008). This tells us the computational speed-up of QDMC. We calculate speed-up for our QDMC simulations with a target time of 2.4 Gyr (green), our QDMC simulations with a target time of 2.8 Gyr (red), and our QDMC simulations with a target time of 3.2 Gyr (blue). We plot the speed-up for both optimistic (dark) and pessimistic (light) QDMC error estimates (Section 2.3). The speed-up is larger for earlier times (when the event is rarer) and reaches a value of about 100 for optimistic error bars.

4. Discussion

Our direct numerical simulations as well as those of Zeebe (2015a) and Laskar & Gastineau (2009) suggest that the probability of Mercury’s orbit becoming unstable in the next 3 Gyr is $\mathcal{O}(10^{-3})$ and in the next 5 Gyr is $\mathcal{O}(10^{-2})$ (Figure 3). Our QDMC results indicate that the probability of Mercury’s orbit becoming unstable in the next 2 Gyr is $\mathcal{O}(10^{-4})$ (Figure 5). We can use these data to speculate that the rate of increase of the log of the probability of Mercury’s orbit becoming unstable decreases with time. This would suggest that the probability of Mercury’s orbit becoming unstable in the next 1 Gyr is less than $\mathcal{O}(10^{-5})$. This speculation could be checked with QDMC simulations with a target time of 1 Gyr in future work, although it is possible this would require some modification of the reaction coordinate and/or more simulations than we performed here.

We performed a direct numerical simulation suite before applying QDMC (Section 2.2). This was necessary in order to confirm that our modeling choices yielded results consistent with previous efforts (Section 3.1) and to confirm that the QDMC results were consistent with the direct numerical simulation results where they overlapped (Section 3.2). Additionally, the direct numerical simulation suite was useful for developing a good reaction coordinate (Section 2.4) since this was the first effort to apply QDMC to a planetary dynamics problem. It is important to note, however, that it is not necessary to perform a direct numerical simulation suite in order to apply QDMC. If you have a good reaction coordinate in hand, you can apply QDMC from the start. Moreover, if you have a reasonable guess at a reaction coordinate, you can run QDMC using it, then use the output to find an improved reaction coordinate and repeat if necessary.

Our estimate of the speed-up associated with QDMC may seem optimistic given that we had to devote a significant

amount of time and effort to finding a good reaction coordinate. It is true that such effort is necessary when a completely new problem is approached using QDMC, but it is not necessary when applying QDMC again to a similar problem. Our speed-up estimates therefore give a reasonable expectation for others hoping to apply QDMC to a related planetary dynamics problem. It is also important to note that, given our computational resources, we would not have been able to estimate the Mercury instability event statistics as accurately as we did using direct numerical simulation. For problems like this, applying QDMC is essential even if non-negligible up-front work is required.

We have demonstrated the utility of QDMC here on an efficient N -body scheme that may not be accurate when Mercury’s eccentricity becomes large (Zeebe 2015b). In future work it would be possible to apply QDMC to the more expensive numerical scheme used by Zeebe (2015b). Alternatively, QDMC could be applied to a scheme like that used in this paper, and then the scheme of Zeebe (2015b) could be switched to once Mercury’s eccentricity became large.

As an alternative to QDMC, there may be astronomical problems for which extreme value theory (De Haan et al. 2006) provides an efficient approach for computing rare-event probabilities. However, extreme value theory rests on statistical assumptions that are not always satisfied (Naess & Gaidai 2008). The assumptions of stationarity and of a mean upcrossing function of generalized Gumbel form are not valid for Mercury’s eccentricity data (Figure 2), and more generally the assumptions behind extreme value theory are invalid or difficult to validate for many problems in the Earth and planetary sciences (Stein 2020). In contrast, QDMC provides unbiased rare-event statistics for any stochastic simulation model, making it a broadly attractive approach for computing rare probabilities.

QDMC could be used to assess the probability that a near-Earth asteroid impacts Earth, which could improve estimates of their value on the Torino Impact Hazard Scale (Binzel 2000). Additionally, a different rare-event scheme such as action minimization (E et al. 2004; Plotkin et al. 2019; Woillez & Bouchet 2020) could be used to calculate the minimum impulse necessary to prevent an asteroid from colliding with Earth as long as a numerical scheme with a linear adjoint is available.

An overdensity is apparent when Mercury’s eccentricity is ~ 0.32 in Figure 2. This overdensity is apparent for running means of $\mathcal{O}(10^7)$ yr, but not in the high-frequency output. The overdensity could be related to the overlap of the g_5 and s_2 resonance modes, which occurs at roughly the same eccentricity of Mercury (Lithwick & Wu 2014). The explanation for the overdensity might be an interesting problem for future research.

In Section 2.4, we demonstrated that the range of values attained by Mercury and Venus’s MOID over a timescale of 400 Myr is strongly correlated with the probability that Mercury will undergo an instability over the ensuing 400 Myr. We found that this reaction coordinate was successful for our purposes since QDMC achieves a large speed-up as compared to direct numerical simulation. However, it is possible that we could improve the efficiency further by finding an improved reaction coordinate, especially using the large sample of simulations we produced using QDMC in which Mercury’s orbit becomes unstable. The underlying dynamical mechanisms leading to Mercury’s potential

instability have been explored by a number of authors (Lithwick & Wu 2011; Boué et al. 2012; Lithwick & Wu 2014; Batygin et al. 2015). This work has established the role of various secular resonances in driving chaotic evolution and Mercury’s eventual eccentricity excitation (see also Laskar 1990; Laskar et al. 1992; Sussman & Wisdom 1992). A reaction coordinate that measures the influence of these resonances, e.g., through frequency analysis of the system’s secular modes (Mogavero & Laskar 2021), could potentially provide improved algorithm performance. More generally, improving the reaction coordinate is a rich problem for future research with many potential angles of attack.

Over extremely short time horizons (<1 Gyr), Mercury’s instability probability may depend sensitively on the details of the simulation model. Unresolved physics (e.g., asteroids, satellites of planets, Pluto, Trans-Neptunian objects, Oort Cloud objects, and other stars passing near the Solar System), the choice of integrator time step, and the details of the noise model may have a significant impact on calculated instability probabilities. Even rounding errors can cause a measurable deviation in orbital paths over such short timescales. However, our working hypothesis, and that of other authors in the field (Laskar 1994, 2008; Batygin & Laughlin 2008; Laskar & Gastineau 2009; Zeebe 2015a; Woillez & Bouchet 2017), is that Mercury instability events on time horizons of 2.4, 2.8, and 3.2 Gyr are not so unlikely that these model details are essential. Testing this hypothesis on shorter time horizons is an outstanding but computationally expensive task.

5. Conclusion

We have shown that QDMC provides unbiased estimates of the probability of rare events in an N -body planetary dynamics code at greatly reduced computational cost. More specifically, using QDMC we were able to estimate the probability that Mercury’s orbit becomes unstable ~ 2 billion years in the future as $\mathcal{O}(10^{-4})$ at a computational cost up to 100 times lower than direct numerical simulation. QDMC is easy to implement and could be applied to many planetary dynamics problems involving rare events.

We thank Edwin Kite for suggesting we investigate Mercury instability events using QDMC. We thank Mark Hammond for help setting up a Conda environment and iPython notebook. We thank Garrett Brown and Hanno Rein for specifying parameter choices for the Rebound simulations. We thank Richard Zeebe for sharing his simulation data with us. We thank Nora Bailey, Jade Checlair, Dan Fabrycky, Gregory Gilbert, Ben Hayworth, Xuan Ji, Tad Komacek, Darryl Seligman, Eric Vanden-Eijnden, and Huanzhou Yang for discussions and comments. We thank Nora Bailey, Konstantin Batygin, Freddy Bouchet, Garrett Brown, Dan Fabrycky, Jacques Laskar, Hanno Rein, Richard Zeebe, and an anonymous reviewer for feedback on an early version of this paper. This work was completed with resources provided by the University of Chicago Research Computing Center. This work was supported by the NASA Astrobiology Program grant No. 80NSSC18K0829 and benefited from participation in the NASA Nexus for Exoplanet Systems Science research coordination network. RJW was supported by New York University’s Dean’s Dissertation Fellowship and by the National Science Foundation through award DMS-1646339. JW acknowledges support from the Advanced Scientific

Computing Research Program within the DOE Office of Science through award DE-SC0020427. SH gratefully acknowledges the CfA Fellowship.

Appendix Implementation Details for QDMC

Here, we explain in more detail our procedure for randomly choosing the number of children N_i^j for a simulation X_i^j and assigning weights to the children. For simplicity, we assume that the unstable trajectories have been set aside for analysis later, and the indices are sorted so that $i = 1$ corresponds to the lowest reaction coordinate value and $i = i_{\max}$ corresponds to the highest reaction coordinate value. Then, we perform the following three steps.

Step one. We partition the indices $i = 1, \dots, i_{\max}$ into sets

$$S^1 = \{L^1, L^1 + 1, \dots, U^1\}, \quad (\text{A1})$$

$$S^2 = \{L^2, L^2 + 1, \dots, U^2\}, \quad (\text{A2})$$

$$\dots, \quad (\text{A3})$$

$$S^{\max} = \{L^{\max}, L^{\max} + 1, \dots, U^{\max}\}, \quad (\text{A4})$$

where the indices (L^i, U^i) are determined by setting $U^0 = 0$ and setting

$$L^i = U^{i-1} + 1, \quad U^i = \min \left\{ j \geq L^i : \sum_{k=L^i}^j \mathbb{E}N_i^k \geq 1 \right\}, \\ i = 1, 2, \dots, i_{\max}. \quad (\text{A5})$$

Step two. For each set S^i , we choose a random number

$$N_i(S^i) = \begin{cases} \left\lfloor \sum_{j \in S^i} \mathbb{E}N_i^j \right\rfloor + 1, & \text{with probability } \sum_{j \in S^i} \mathbb{E}N_i^j - \left\lfloor \sum_{j \in S^i} \mathbb{E}N_i^j \right\rfloor, \\ \left\lfloor \sum_{j \in S^i} \mathbb{E}N_i^j \right\rfloor, & \text{otherwise.} \end{cases} \quad (\text{A6})$$

Instead of choosing the numbers $N_i(S^i)$ independently, we use pivotal sampling (Deville & Tille 1998) to ensure the random numbers $N_i(S^i)$ satisfy (A6) and also $\sum_{i=1}^{i_{\max}} N_i(S^i) = N$.




Step three. For each index $j \in S^i$, we choose a random number

$$N_i^j = \begin{cases} \left\lfloor \frac{N_i(S^i)w_i^j}{\sum_{k \in S^i} w_i^k} \right\rfloor + 1, & \text{with probability } \frac{N_i(S^i)w_i^j}{\sum_{k \in S^i} w_i^k} - \left\lfloor \frac{N_i(S^i)w_i^j}{\sum_{k \in S^i} w_i^k} \right\rfloor, \\ \left\lfloor \frac{N_i(S^i)w_i^j}{\sum_{k \in S^i} w_i^k} \right\rfloor, & \text{otherwise,} \end{cases} \quad (\text{A7})$$

We use pivotal sampling to ensure the numbers N_i^j satisfy (A7) and also $\sum_{k \in S^i} N_i^k = N_i(S^i)$, and we assign the children of X_i^j updated weights $\sum_{k \in S^i} w_i^k / N_i(S^i)$.

For more discussion motivating these choices, see Webber et al. (2020) and Greene et al. (2021).

ORCID iDs

Dorian S. Abbot  <https://orcid.org/0000-0001-8335-6560>
Robert J. Webber  <https://orcid.org/0000-0001-8286-6315>
Sam Hadden  <https://orcid.org/0000-0002-1032-0783>

References

- Allen, R. J., Frenkel, D., & Rein ten Wolde, P. 2006, *JChPh*, **124**, 024102
- Antoszewski, A., Feng, C.-J., Vani, B. P., et al. 2020, *JPCB*, **124**, 5571
- Batygin, K., & Laughlin, G. 2008, *ApJ*, **683**, 1207
- Batygin, K., Morbidelli, A., & Holman, M. J. 2015, *ApJ*, **799**, 120
- Binzel, R. P. 2000, *P&SS*, **48**, 297
- Boué, G., Laskar, J., & Farago, F. 2012, *A&A*, **548**, A43
- Brown, G., & Rein, H. 2020, *RNAAS*, **4**, 221
- Cérou, F., & Guyader, A. 2007, *Stoch. Anal. Appl.*, **25**, 417
- De Haan, L., Ferreira, A., & Ferreira, A. 2006, *Extreme Value Theory: an Introduction*, Vol. 21 (New York: Springer)
- Dean, T., & Dupuis, P. 2009, *Stoch. Anal. Appl.*, **119**, 562587
- Dean, T., & Dupuis, P. 2011, *Ann. Oper. Res.*, **189**, 63
- Deville, J.-C., & Tille, Y. 1998, *Biometrika*, **85**, 89
- Dupuis, P., & Wang, H. 2004, *Stochastics*, **76**, 481
- E, W, Ren, W., & Vanden-Eijnden, E. 2004, *Commun. Pure Appl. Math.*, **57**, 637
- Giardina, C., Kurchan, J., & Peliti, L. 2006, *PhRvL*, **96**, 120603
- Grassberger, P. 1997, *PhRvE*, **56**, 3682
- Greene, S. M., Webber, R. J., Berkelbach, T. C., & Weare, J. 2021, arXiv:2103.12109
- Guttenberg, N., Dinner, A. R., & Weare, J. 2012, *JChPh*, **136**, 06B609
- Haraszti, Z., & Townsend, J. K. 1999, *ACM Trans. Model. Comput. Simul.*, **9**, 105
- Huber, G., & Kim, S. 1996, *BpJ*, **70**, 97
- Kahn, H., & Harris, T. 1951, *Natl. Bureau Stand. Appl. Math. Ser.*, **12**, 27
- Laskar, J. 1989, *Natur*, **338**, 237
- Laskar, J. 1990, *Icar*, **88**, 266
- Laskar, J. 1994, *A&A*, **287**, L9
- Laskar, J. 2008, *Icar*, **196**, 1
- Laskar, J., & Gastineau, M. 2009, *Natur*, **459**, 817
- Laskar, J., Quinn, T., & Tremaine, S. 1992, *Icar*, **95**, 148
- Lithwick, Y., & Wu, Y. 2011, *ApJ*, **739**, 31
- Lithwick, Y., & Wu, Y. 2014, *PNAS*, **111**, 12610
- Liu, J. S. 2008, *Monte Carlo Strategies in Scientific Computing*, Vol. 10 (New York: Springer), 978
- Mogavero, F., & Laskar, J. 2021, *A&A*, **655**, A1
- Moral, P. D., & Garnier, J. 2005, *Ann. Appl. Probab.*, **15**, 2496
- Naess, A., & Gaidai, O. 2008, *J. Eng. Mech.*, **134**, 628
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825
- Plotkin, D. A., Webber, R. J., O'Neill, M. E., Weare, J., & Abbot, D. S. 2019, *JAMES*, **11**, 863
- Ragone, F., & Bouchet, F. 2021, *GeoRL*, **48**, e2020GL091197
- Ragone, F., Wouters, J., & Bouchet, F. 2018, *PNAS*, **115**, 24
- Rein, H., & Liu, S.-F. 2012, *A&A*, **537**, A128
- Rein, H., Tamayo, D., & Brown, G. 2019, *MNRAS*, **489**, 4632
- Rosenbluth, M., & Rosenbluth, A. 1955, *JChPh*, **23**, 356
- Stein, M. L. 2020, *StaSc*, **35**, 31
- Sussman, G. J., & Wisdom, J. 1992, *Sci*, **257**, 56
- Tamayo, D., Rein, H., Shi, P., & Hernandez, D. M. 2020, *MNRAS*, **491**, 2885
- Tibshirani, R. 1996, *J. R. Stat. Soc. Series. B. Stat. Methodol.*, **58**, 267
- Vanden-Eijnden, E., & Weare, J. 2012, *Commun. Pure Appl. Math.*, **65**, 1770
- Villen-Altamirano, M., & Villen-Altamirano, J. 1991, in *Proc. of the 13th Int. Teletraffic Congress, Queuing, Performance and Control in ATM*, ed. J. W. Cohen & C. D. Pack (Amsterdam: North-Holland), 71
- Warmflash, A., Bhimalapuram, P., & Dinner, A. R. 2007, *JChPh*, **127**, 154112
- Webber, R. J. 2022, *European Nonlinear Dynamics Conf. 2022 (Lyon, France)*
- Webber, R. J., Aristoff, D., & Simpson, G. 2020, arXiv:2011.13899
- Webber, R. J., Plotkin, D. A., O'Neill, M. E., Abbot, D. S., & Weare, J. 2019, *Chaos*, **29**, 053109
- Wisdom, J. 2006, *AJ*, **131**, 2294
- Wisdom, J. 2015, *AJ*, **150**, 127
- Wisdom, J., & Holman, M. 1991, *AJ*, **102**, 1528
- Wisdom, J., Holman, M., & Touma, J. 1996, *Fields Inst. Comm.*, **10**, 217
- Wuillez, E., & Bouchet, F. 2017, *A&A*, **607**, A62
- Wuillez, E., & Bouchet, F. 2020, *PhRvL*, **125**, 021101
- Wolf, E., & Toon, O. 2015, *JGRD*, **120**, 5775
- Zeebe, R. E. 2015a, *ApJ*, **811**, 9
- Zeebe, R. E. 2015b, *ApJ*, **798**, 8