Chapter 6

Modelling the Complete Problem

This chapter attempts to construct a self-consistent model for launch and delivery of the lunar and martian meteorites. The major problem for this enterprise is we have a poor understanding of the launch physics. For this reason, this thesis has tried very hard to not invoke any particular model of cratering physics. Up to this point, none of the results have depended on any model of the ejection physics; all that has been important so far is the celestial mechanics of the particles *after* launch (and some collisional disruption in the case of martian meteoroids). Of the five puzzles regarding the lunar and martian meteorites (p. 7), only that regarding the fairly short 4π ages of the martian meteorites has been resolved (Chpt. 4). The most puzzling of the remaining facts is the approximately equal numbers of lunar and martian meteorites. My chief goal is to understand how to reduce the expected ratio of lunar to martian meteorites from order 10^3 (Wetherill 1984) to order 1. I begin by reviewing the constraints available from the meteorite data, the simulations, and proposed theories for cratering ejection at high velocity.

6.1 Constraints

6.1.1 Meteorite data

The greater average masses of the martian meteorites has already been noted. Note that the four martian falls (Table 2.1) account for four of the five large (>3 kg) martian meteorites; this could be blamed on a selection effect, but one must still explain why there are no comparably large lunar falls. The depth data show that all the lunar meteorites with short blast-off times ($t_{bo} < 0.3$ Myr; p. 15) originate at shallow depths. Only the three meteorites that left the Moon more than ~1 Myr ago were buried deeply at the time of their ejection. In contrast, all the martian meteorites were buried very deeply at the time of their ejection.

6.1.2 Source-crater pairing (SCP)

Warren (1994) concluded that at least six source craters were required to launch the (then nine) lunar meteorites. More recently, Thalmann *et al.* (1996) claim at least six, and possibly eight, separate craters for the eleven meteorites listed in Table 2.1. Both works feel that a source-crater pairing (SCP) between Y-793169 and Asuka 881757 is very probable; the recently determined difference in the depth provenance of these two meteorites (Thalmann *et al.* 1996) probably rules out their being launched in a single meteoroid and paired in the conventional sense. In addition, Warren suggests as plausible a SCP between EET87521 and Y-793274, but rejects claims for a SCP between ALH81005 and Y-791197, as well as any pairing of these latter two with any others. Source-crater pairings, although interesting, do not affect comparisons of the meteorite data with the simulated CRE age spectra, since the latter are only sensitive to the fraction of transits that occur in a given time, and SCP bodies are delivered independently. I believe that the overlap of (relatively uncertain) cosmic-ray exposure ages is a necessary, but hardly sufficient, condition to establish SCPs; the petrologic data are at least equally important. Thus if isotopic ratios (Warren 1994, Thalmann *et al.* 1996) in two meteorites that are SCP candidates do not match, the pair can be rejected since such ratios are known to vary little over the length scale of Apollo traverses.

Martian meteorites separate rather naturally into age *and* petrographic groups. ALH84001 is distinct in both ways. The three nakhlites were all launched about 10.5 Myr ago (again, errors of 20–30% in these ages are likely). Chassigny, at $t_{bo} \simeq 9.9$ Myr (Eugster *et al.* 1996), cannot be split from the nakhlites on CRE grounds, but its petrology is so distinct that it conceivably represents a separate event; since three of the NCs are isolated falls, none of these four can be paired in the conventional sense. Eugster *et al.* (1996) provide support that the lherzolitic shergottites CRE ages of $\simeq 3.5$ Myr are distinct from the basaltic shergottites at 0.6 and 2.7 Myr. If these three CRE ages are due to three martian impact ejection events, then it is rather curious that no other shergottite ejection events are present in the record after 4 Myr (since presumably this implies that such events are common); it is also puzzling then that no events *other* than shergottite ejections are represented during the last 10 Myr. If all the shergottites were spawned in a single impact (presumably ~ 3.5 Myr ago), then this problem is reduced, but in this case the launch of large bodies is required to shield the younger meteorites, with all the attendant difficulties of this hypothesis (Sec. 4.4).

6.1.3 Source regions on Mars

Warren (1994) points out that the lunar meteorites seem to representatively sample the various geological regions of the lunar surface, and thus we have no reason to speculate that any region of the Moon is particularly favored by the launch physics. The simulations of Chpt. 3 also indicate that no regions of the lunar surface are dynamically favored for delivery. In contrast, as mentioned in Chapters 1 and 2, the martian meteorites are distinguished (except for ALH84001) by the fact that they all have young crystallization ages. This prevalence of young ages (11 of the 12 meteorites) is unexpected since only a small fraction ($\simeq$15%) of the martian surface is thought to be this young (Tremain 1995a). This fact favors a small number of launch events (even one), for then the youth of the samples could be attributed to small-number statistics. If one postulates only 3 launch events (one each for S, NC, and ALH84001), then the probability that 2 of them occur in young terrain is $\approx$5%, which is perhaps acceptable; however, the 10-Myr impact must overlap both N and C terrain as well as launch large bodies in order to produce different shielding times for the shergottites (as just discussed). If more than 3 impacts are responsible for launching the martian meteorites, then sampling only young terrain by chance becomes very unlikely. Treiman (1995a) presents a more detailed discussion of this topic.

Warren (1994) has proposed that the lack of geologically old martian rocks is not due to chance, but instead indicates something important about the surface properties of the older regions of Mars. He proposes that weathering processes have degraded the surface structure such that meteorite launch from older portions of the surface is inhibited. This hypothesis would then offer the following explanations: (1) we collect fewer martian than lunar pebble-sized meteorites because coherent un-weathered pebbles are rare on Mars; (2) few martian meteorites are launched from shallow depths because only a small fraction of the outer several meters of Mars have escaped weathering; and (3) SNC meteorites are young mafic rocks because (a) only young rocks are free of pervasive weathering; (b) even without weathering, sedimentary rocks and most impact breccias are not strong enough to survive the rigors of ejection; and (c) most young, igneous rocks on Mars are probably mafic.

Sec. 6.3.3.2 discusses the influence on the transfer model of restricted portions of the martian surface being available for launch, including the possibility that the martian atmosphere favors launch from terrains at high altitude. It is likely that the meteorite data will never be able to answer this question definitively; only a sample return mission to Mars will.

6.1.4 The simulations

We have already seen that the time scales for lunar and martian meteorite delivery agree well with the 4π ages of the meteorites. In the case of the Moon, the prevalence of quick transfers in the time spectrum of the deliveries leads one to conclude that most of the ejecta barely escapes the Moon, *i.e.,* the ejecta's velocity spectrum must be relatively steep; this will be quantified below. In contrast, the ejection velocity makes little difference to the CRE age spectrum from Mars, so we are unable to constrain this using the 4π ages of the martian meteorites, only the average collisional lifetimes (relatively unimportant for the lunar case) of martian meteoroids were deduced.

In both cases the simulations provide the time spectra for delivery as a function of launch velocity, and the absolute delivery efficiency. The latter yield is approximately Y=40±15% for the Moon, depending on the launch velocity, and Y=4±2% from Mars. The error in the martian delivery efficiency is due to small number statistics in the simulations (only 10-20 particles from each 300-particle simulation impacted Earth) and the unknown collisional effects (both the time scale and whether all collisions should be viewed as catastrophic). The roughly 50 percent uncertainty in these absolute delivery efficiencies will be more than sufficient for our purposes, since we are only trying to match fluxes to order of magnitude.

It is useful at this point to define two quantities. If $n_\oplus(t)dt$ is the number of (lunar or martian) meteorites arriving at Earth between t and $t + dt$, then the Earth transfer efficiency $\mathcal{E}(t)$ is given by

$$n_\oplus(t)\ dt = N_{esc}\,\mathcal{E}(t)\ dt\ , \tag{6.1}$$

where N_{esc} is the number of objects that escaped from the parent body in some impact event. (Note that, for clarity, I always use the term 'impact' to refer to the launch event; the term 'arrival' or 'delivery' is used for the reception of the meteorite at Earth). Then for any given impact, the total number of objects that have ever arrived on Earth by time T after launch is

$$N_\oplus(T) = N_{esc}\int_0^T \mathcal{E}(t)\ dt \equiv N_{esc}\times Y\times\mathcal{F}(T)\ \ , \tag{6.2}$$

where Y is the total yield (the fraction of launched particles that will ever reach Earth) given above, and $\mathcal{F}$ is the cumulative fraction of material arrived. In general, $\mathcal{F}(T)$ satisfies $\mathcal{F}(0) = 0$ and $\mathcal{F}(\infty) = 1$, and is basically given by the cumulative CRE age plots presented earlier (Figs. 3.13 and 4.7). Of course, since the simulations are of finite duration they only provide an approximation to $\mathcal{F}(T)$. Since the delivery rate has dropped so much in the later part of the simulations and Sun-grazing or collisional destruction will be destroying ever-increasing numbers of meteoroids, I believe that relatively few particles will arrive at Earth after the termination of the simulations and thus the simulated spectra are sufficient.

Note that the cumulative plots do not indicate the differences in total yield Y between the simulations, which can be obtained from Tables 3.5 and 4.2. For the Moon $\mathcal{F}$(1 Myr) > 0.8 for all the simulations, so that almost all of the lunar meteorites that will arrive at Earth will do so in <1 Myr. For the 2.4 km/s simulation, the delivery is so rapid that $\mathcal{F}$(0.7 Myr) $=$ 0.9, *i.e.,* only 10% of meteorites take longer than 0.7 Myr to arrive at Earth. The lunar and martian

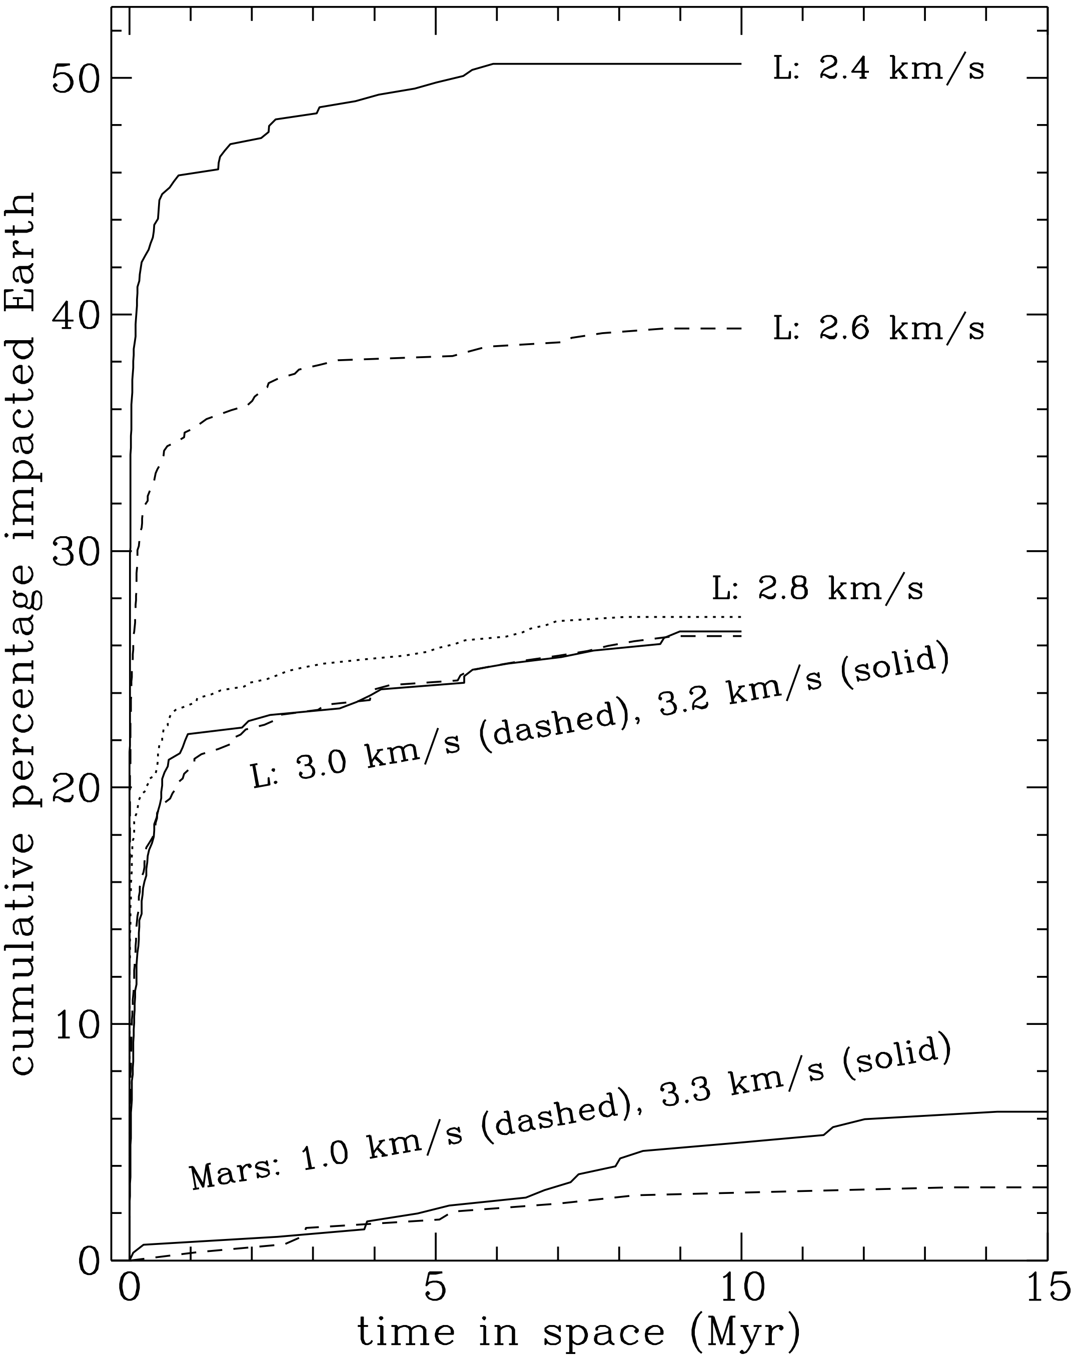

Figure 6.1 The percentage of *launched* particles that impact Earth from lunar (L) and martian simulations. The lunar curves are marked with the ejection velocities, while the martian speeds are v_∞'s. The vertical axis is $Y \times \mathcal{F}(T)$ (*cf.* Eq. 6.2).

$N_\oplus(T)$ curves are plotted using a linear time axis in Fig. 6.1, which brings home how sharply the delivery of lunar material tails off compared to that from Mars.

6.1.4.1 Age of the ice sheet

It is important to realize that the bulk of our meteorite sample was collected in Antarctica (Sec. 1.2). The foregoing discussion does not take into account that the ice sheet is of finite age (denoted $A_{\rm ice}$). This age should not be viewed strictly as how old the ice is, but rather the typical lifetime of a meteorite in the ice before it is destroyed.

Eq. 6.2 gives the number of meteorites that have arrived on Earth since the lunar or martian impact T years ago. However, if the launch occurred longer ago than $A_{\rm ice}$, then many of the meteorites that landed on the ice sheet have since been destroyed. Therefore, the actual number of meteorites currently in the ice sheet is given by

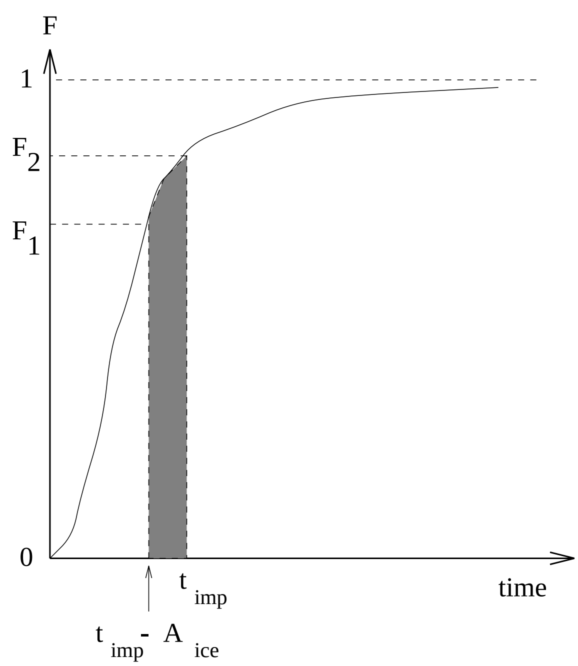

Figure 6.2 The ice sheet will only accumulate meteorites during its age $A_{\rm ice}$, and thus only the fraction $\mathcal{F}_2 - \mathcal{F}_1$ of all the meteorites that arrive are currently in the ice sheet if the impact occurred before $A_{\rm ice}$.

$$N_\oplus = N_{esc} \times Y \times \begin{cases} \mathcal{F}(t_{imp}) & \text{if } t_{imp} < A_{\rm ice} \\ \mathcal{F}(t_{imp}) - \mathcal{F}(t_{imp} - A_{\rm ice}) & \text{if } t_{imp} > A_{\rm ice}\ , \end{cases} \tag{6.3}$$

where t_{imp} is how long ago the impact occurred. This is illustrated in Fig. 6.2. The effect will be especially important for the lunar meteorites since most of them arrive in a short burst after the impact. If this burst finishes more than A_{ice} ago, then most of the meteorites from that impact are no longer recoverable.

The age of the ice sheet can be estimated from studies of the CRE data for Antarctic meteorites (Cassidy and Rancitelli 1982; Nishiizumi *et al.* 1989). The number of Antarctic meteorites drops almost exponentially with terrestrial age ($t_{\oplus}$), with a half-life of perhaps just slightly less than 100 kyr (Nishiizumi *et al.* 1989). The *median* terrestrial age of the lunar meteorites is somewhere between 30 kyr to 60 kyr; unfortunately many of the $t_{\oplus}$'s are just upper limits (Table 2.1). The *mean* age may not be a good representation as the 230 kyr age of MAC88105 unrepresentatively skews the distribution. The ice fields, and even the various stranding areas of each ice field, have different distributions of terrestrial ages, so deciding upon a single number may be difficult (and/or misleading). Certainly the 'mean age' is more than 10 kyr and less than 150 kyr. The martian meteorites have only 4 Antarctic terrestrial ages determined: 10, 11, <100, and 200 kyr, and so are not clearly any different in this regard.

The model used in Eq. 6.3, in which the single cutoff $t_{\oplus} = A_{\rm ice}$ is imposed, is perhaps overly simplistic. A more realistic model could include an exponential loss process, in which each delivery has some probability of being destroyed before recovery. We should note that since the average mass of lunar meteorites is smaller than the martian meteorites (or the bulk sample of all meteorites) one might suspect the average retention time to be somewhat smaller. I have not included these additional complications in the current model.

6.1.5 Cratering mechanics

In order to make some estimates of the efficiency of meteoroid launch from an impact so that a complete picture can be developed, at this point we must examine a model for the ejection physics. The results available in the literature are based on hydrocodes (*e.g.,* O'Keefe and Ahrens 1977; Holsapple and Choe 1988), analytical models (*e.g.,* Zel'dovich and Raizer 1967, Melosh 1984), and laboratory cratering experiments (*e.g.,* Gault *et al.* 1963, Gratz 1993, Kato *et al.* 1995).

It is crucial to realize that much of the work that has been done on understanding impact cratering is concerned with the excavation of the crater itself. The crater's cavity is formed by a subsonic excavation flow (see Melosh 1989 for a general introduction to cratering phenomena). In contrast, the physics that ejects the very fastest, lightly shocked material represented by the lunar and martian meteorites is probably quite different from that taking place in the excavation. The ejecta blocks that create secondary crater fields are interpreted to have been launched from the original crater at up to 1 km/s (Vickery 1987), in a process which ejects material from a near-surface spall zone (Melosh 1989). It is plausible, but not proven, that this same process is responsible for the ejection of material at faster than the escape velocity; there is some indication in the work of Gault *et al.* (1963) that some transition occurs in the ejection physics at about 1 km/s (at least for basalt targets).

Any model of the launch process must explain the lightly shocked state of the lunar and martian meteorites (Chap. 1). This historically surprising fact is now viewed as a major constraint on the launch process. Two processes have been suggested in the literature. The first is oblique-impact-induced, vapor-jet plume entrainment (Nyquist 1983, O'Keefe and Ahrens 1986), in which the ejecta are entrained in an expanding gas cloud which accelerates them to escape velocity. The second model relies on the interference of shock waves (Melosh 1984, 1985) in a near-surface spallation layer to propel particles to speeds greater than escape while maintaining low peak pressures. I have chosen to base all of my discussion on the Melosh model because (1) production of lightly shocked, high-velocity spalled ejecta has been observed in laboratory experiments (Gratz 1993), and (2) the quantitative predictions of the model are in reasonable concordance with a number of features of the meteorite data. It is important to realize though that the transfer dynamics are probably not sufficient in themselves to distinguish between these launch models.

6.1.5.1 The spallation model

I will briefly summarize the shock-wave interference model, for which the best introduction is Melosh (1985). Consider a projectile of diameter D striking a surface with a component of velocity U normal to the surface. The impact produces a shock wave which appears to travel spherically outward from the *equivalent center*, a point roughly one projectile diameter below the target's surface. As the shock wave sweeps past any given point, the pressure increases with a characteristic 'rise time' τ equal to the time it took the impactor to penetrate into the surface by its own radius,

$$\tau \sim \frac{D}{2U} \, . \tag{6.4}$$

After this, the pulse decays away in a time $\beta\tau \sim D/c_L$, where c_L is the sound speed in the target, and β is a constant. The total time duration of the pressure pulse is thus $(1+\beta)\tau$. The model relies on the fact that, near the surface of the target, two waves arrive at any given point from the equivalent center: a direct wave, and one reflected from the nearby surface. Since the sign of the reflected pressure wave is reversed, these two waves interfere and in fact the peak pressure reached is considerably smaller than the maximum pressure in the wave. However, even though rocks at the interference zone do not feel the full stress pressure, they are still accelerated because the final velocity of material after the wave has passed depends on the pressure gradient rather than the pressure itself. In the interference zone the pressure actually rises twice as rapidly as the case when no interference occurs, due to the presence of the free surface; this can accelerate material to nearly twice the particle velocity of the wave itself. This mechanism can therefore eject lightly-shocked particles at high speeds.

Melosh's estimate of the depth z of the interference layer is important for our purposes. This layer is a lens-shaped region, the depth of which increases with the horizontal distance s from the impact point:

$$z = \frac{c_L \tau}{2} \sqrt{1 + \frac{s^2}{D^2 - \left(\frac{c_L \tau}{2}\right)^2}} \,. \tag{6.5}$$

From Eq. 6.4, $c_L \tau \simeq c_L D/(2U)$, and so

$$z = \frac{c_L D}{4U} \sqrt{1 + \frac{s^2/D^2}{1 - \left(\frac{c_L}{4U}\right)^2}} \tag{6.6}$$

For a hypervelocity impact ($U > c_L$) we can approximate the denominator of the second term as unity. Since the launch velocity drops off roughly as s^{-3} (Warren 1994), the majority of the highest velocity ejecta should be launched within a few impactor diameters of the contact point ($s/D \sim 2$–3), and thus it is roughly correct to say that the spall zone has an approximately constant thickness of order

$$z \sim D \frac{c_L}{U}. \tag{6.7}$$

For U=10–20 km/s impacts into pristine rock (where c_L ~6 km/s), the spall zone depth is of order an impactor radius. However, impacts into a regolith layer (where $c_L \ll$1 km/s) will have interference zones very much shallower than the impactor. This is of great importance, since the sound speed in the upper few meters of the lunar regolith is measured by the Apollo experiments to be only $c_L \approx$100 m/s (McKay *et al.* 1991). Thus, a 100-m projectile hitting the lunar regolith at 10 km/s engenders an interference zone of depth only one meter! Since the lunar regolith is generally 4–5 m in the mare and perhaps 10–15 m in the highlands (McKay *et al.* 1991), only impactors of order 1 km can mine down below the regolith to launch particles from depth. However, the Melosh spallation model may need to be modified to cope with a layered target (Warren 1994).

Therefore, the low sound speed of the Moon's surface results in a drop in the number of expected meteoroids launched in any single impact, by up to two orders of magnitude. It is not certain that the martian surface should be any more favorable (in a relative sense) for meteorite launch, but the megaregolithic structure (McKay *et al.* 1991) of the lunar surface is probably not present on Mars, especially in the younger terrains, due to ongoing surface processes and recent volcanism on Mars (Christensen and Moore 1992). Therefore one expects a relative dearth of ejections from the Moon, which may be a major reason for the 'over-representation' of martian meteorites in the collection. I will return to this topic below.

6.1.6 Size distributions of surface material

Since a regolith cannot support tensile stress, the surface layer is simply ejected in a spall process (Melosh 1985). In contrast, in a solid target ejecta is spalled off in plates that then break up in flight into fragments whose size can be estimated (Melosh 1989). The 2π CRE evidence tells us that the martian meteorites have come from depths below several meters, and so the size distribution of martian surface material may be irrelevant for understanding the sizes of its meteorites. Nevertheless, for completeness, note that the surface size distribution of martian rocks has been estimated from Viking lander photography by Golombek and Rapp (1996).

In the case of lunar launch, the size distribution of the escaping meteoroids should presumably be that of the regolith since the surface layer has little strength. The size distribution of lunar soils is well catalogued (Graf 1993), but little information is readily available about the relative number of rocks in the range from 1 mm to 10 cm (the size of the meteorites) in the upper few meters of the regolith. Particles in these size ranges are too large to be well represented in drill cores, and it is difficult to reach firm conclusions from radar studies due to model uncertainties (Carrier *et al.* 1991).

The surface photography performed by the Surveyor missions (Shoemaker *et al.* 1969) is the best source of this size information. All of these spacecraft landed in mare areas, except Surveyor VII which landed in highland terrain and next to a large young crater (Tycho). From television pictures of the Surveyor I, III, V, VI, and VII landing sites, the cumulative number of particles N per 100 m^2 was fairly well represented by a relation of the form $N = Kd^{\gamma}$, with γ ranging from -1.8 to -2.6, for particle diameters d from $\simeq$1–500 mm. The absolute normalization varied from site to site, as some sites were especially rich in larger rocks. The surface density distributions were then used by Shoemaker *et al.* (1969) to estimate the volumetric particle-size-frequency distribution functions. Except for the Surveryor VII site (which was atypical due to being so close to Tycho), these authors find that rocks larger than $d \simeq 1$ cm in diameter (mass $\sim$1 g) occupy only $\simeq 2 \pm 1\%$ by volume of the lunar regolith. Thus, even allowing for some uncertainty, to obtain $\sim$1 10-g lunar meteorite requires launching *at least* two orders of magnitude more material (by volume) since so much of the regolith consists of fine soils.

The lunar meteorite size distribution is statistically consistent with a d^{-2} power-law, but only for those meteorites larger than 30 g. This is not surprising; collected Antarctic meteorites are believed to be representative of those actually in the ice only at masses above this limit (Huss 1990). There is therefore no reason to suspect that the lunar launch process favors any particular mass range.

6.2 Modelling lunar launch

I shall now adopt the goal of using the meteoroid dynamics to determine as much information as possible about the lunar launch process.

6.2.1 Launch rate, lower limit

The observed number of lunar meteorites (LMs) can be used in conjunction with transfer efficiency to estimate a lower limit to the meteoroid launch rate. Consider the following four assumptions:

1. All LMs in the collection have been identified,
2. LMs are no more difficult to recognize than other meteorites,
3. LMs are not more prone to destruction by weathering,
4. LMs are just as strong during entry as other meteorites.

If any of these assumptions are incorrect, then the launch rate we are about to estimate will rise. Now, of the $\simeq$ 20,000 Antarctic meteorites, 15 are lunar if we count all the meteorites uncorrected for pairing, or of order one in a thousand if assumptions 1–3 are correct. If assumption 4 is correct, then this ratio of 1 in 10^3 is also representative of the fraction of meteoroids that are presently in Earth-crossing space (at least, over the last $A_{\rm ice}$ years). Therefore we can use reasonably well calibrated estimates of the total flux of meteorites (Halliday and Griffen 1992; Halliday *et al.* 1984, 1989, 1991; Huss 1990) to estimate the lunar fall rate. The smallest flux estimate in the literature is that of Halliday *et al.* (1989) who estimate the number of incident meteoroids larger than some total meteorite mass m_T (in grams) to be

$$N(m > m_T) = 260 m_T^{-0.49} \quad /10^6 {\rm km}^2/{\rm yr} \tag{6.8}$$

for meteorites with total masses of <1 kg. The smallest lunar meteorite is 6 g, and probably had a pre-atmospheric mass of at least 10 g. Using this 10-g mass as a lower bound in Eq. 6.8, the fall rate of meteorites is $\sim$100 /10^6 km^2 /yr, and therefore the lunar fall rate is at least $\sim$0.1 /10^6 km^2 /yr. Accordingly, the Earth's surface area receives an integrated flux of at least 50 lunar meteorites per year. The lunar transfer efficiency is 40$\pm$15%, and using 50% gives us a lower bound on the (averaged) launch rate of at least 100 meteoroids per year. This launch-rate estimate applies to strong meteoroids bigger than 10 g over the last Myr or so that the CRE ages span. Given that we have had several lunar launch events in the last 10^5 years, the impactors are probably striking the Moon every $\sim 10^4$ years, implying a time-averaged estimate of at least 10^6 meteoroids launched per event.

It is unlikely that any of the 4 assumptions given above are wrong even to a factor of 2. The strength assumption (#4) is the most interesting, as Warren (1994) points out that lunar meteorites tend to have much higher crushing strengths than Apollo samples. It is not clear, however, whether lunar meteorites are systematically stronger or weaker than chondrites, or if crushing strength is the relevant parameter for survivability during atmospheric entry. Nevertheless, let us assume that the lunar meteorites are of comparable strength to other meteorites (or alternately that the launch estimates are only of rocks at least this strong), and calculate a best estimate of the launch rate instead of a lower limit. Huss (1990) believes the most realistic estimate of meteorite infall rate from 1 g to 1000 kg to be

$$N(m > M) = 6800\ M^{-0.833}\ \ /10^6\mathrm{km}^2/\mathrm{yr}\ , \tag{6.9}$$

where M is the terminal mass (in grams) of the meteorite; Halliday *et al.* (1996) give a relationship that is similar within a factor of 2 for masses >2 kg. The factor of a few difference between the terminal and total mass need not concern us much here, but let us use a terminal mass of 6 g for the sake of argument. Assuming that 1 in 10^3 of this flux is lunar yields a best estimate of the lunar launch rate of ~1500 meteoroids/year being injected into near-Earth space (in a time-averaged sense).

6.2.2 Impactor flux

The flux of lunar impactors is estimated by crater counting on the Moon and Earth (Melosh 1989, Neukum and Ivanov 1994, Hartmann 1995) and telescopic observations of near-Earth objects (Shoemaker 1984, Rabinowitz 1994, Shoemaker *et al.* 1994, Bottke *et al.* 1994). Observationally, the impactor population can be described by a *cumulative* distribution N(diameter > D) satisfying:

$$N(D)\,dD = CD^{-\beta}\,dD\ , \tag{6.10}$$

with $\beta = 2$ for $D > 100$m. There is some evidence (Rabinowitz *et al.* 1994, Warren 1994) for a steepening of the size distribution to D^{-3} below 100 m. I will examine below the effect of allowing the exponent to change to $\beta = 3$ at sizes less than 100 m.

I bin the impactors into an incremental size distribution (see Melosh 1989, p. 186) with bins of size $(D, \sqrt{2}D)$, *i.e.,* the upper boundary of each bin is $\sqrt{2}$ of its lower boundary. This binning has the property that the number of impactors per bin also scales as D^{-2} (or -3). That is,

$$N(D, \sqrt{2}D) \propto D^{-\beta} \tag{6.11}$$

I have pegged the impact *rate* to the frequency of a 1-km impactor, and scaled all others via the number distribution above. According to many sources (*e.g.,* Bottke *et al.* 1994), an impactor larger than 1 km currently hits the Moon every $t_{peg} \simeq 2$ Myr, with probably at most a factor of 3 error. This implies that an impactor between 1 and $\sqrt{2}$ km hits the Moon every 4 Myr. Some have suggested, however, that at present the lunar impactor flux is twice that of several Gyr ago (Shoemaker 1983, Moore and McEwen 1996). The lunar CRE age spectrum turns out to be relatively insensitive to these uncertainties, although the absolute flux of meteorites will of course depend linearly on the impactor flux.

I have arbitrarily chosen the upper end of the largest impactor diameter bin to be at $D = 32$ km, thus the last bin is (22.62,32) km. For $\beta = 2$, an impactor with $D > 32$ km currently hits the Moon only every $\sim$1 Gyr, so ignoring events less frequent than this will be irrelevant to the study of meteorites that are now in the Antarctic ice sheet. This upper bin choice also has the convenience of having one bin boundary at $D = 1$ km. The sensitivity of the results to the *lower* end of the impactor size distribution will be discussed. For impacts to occur at least every 10,000 years the smallest impactors must be < 100 m.

6.2.3 Launch scaling

It is now unavoidable that some form of scaling be introduced for the cratering events. I will attempt to make minimal, and reasonable, assumptions. Clearly we expect larger, more energetic impacts to liberate more particles than smaller ones, and thus the launch scaling I will use is that the number $\mathcal{N}$ of particles launched between v and $v + dv$ drops off as a power law, with exponent p:

$$\mathcal{N}(v)dv \propto \Phi(E)v^{-p}dv \ , \tag{6.12}$$

where Φ is some monotonically increasing function of the impactor's kinetic energy E. Power-law relationships like Eq. 6.12 are often observed as the results of cratering experiments (Gault *et al.* 1963, Frisch 1992, Kato *et al.* 1995). This relation implies that the number of *escaping* particles is

$$N_l(p, E, v_e) = \int_{v_{esc}}^{\infty} \mathcal{N}(v)dv = k\frac{\Phi(E)}{p-1}\left(\frac{1}{v_{esc}}\right)^{p-1} \ . \tag{6.13}$$

I am assuming that the impactors all arrive with a median velocity, and a velocity spectrum that is independent of size. The former assumption is known to be wrong. Chyba *et al.* (1994) shows the distribution of lunar impact velocities for Earth-crossing asteroids; the median value of v_∞ is only 12 km/s, the majority occur at under 15 km/s, and 25% of the impacts occur at <10 km/s. Less than 10% of lunar impacts occur at >25 km/s, which would boost the energy by more than a factor of 4 compared with the median value (for an impactor of a given diameter). I will assume that all impactors arrive at the median velocity, or at least that the increased effect of fast impactors of a given size is cancelled out by the underestimated effect of slower speed impactors of a larger size. This single-speed hypothesis will be less valid if comets (which impact at higher velocities and for which the ejection physics may differ) make up a reasonable fraction of the relevant impactors, or if it is only the highest velocity impactors that produce most of the ejecta. In any case, removing the impactor velocity as a parameter means that the impactor's kinetic energy is a function only of the impactor's diameter D, and

$$\Rightarrow \Phi(E) = \Phi(D) \; . \tag{6.14}$$

Following Warren (1994), I adopt a power-law expression for the number of meteoroids launched as a function of the impactor diameter of the form

$$\Phi(D) \propto D^{\alpha} \; . \tag{6.15}$$

If ejected meteoroids were all of comparable diameter, then Vickery and Melosh (1987) calculate $\alpha = 3.84$, but Melosh (private communication, 1996) is cautious about this result since it is based on a 2-D hydrocode. Warren (1994) believes α must be near 2 for the Moon, but that $\alpha \approx 3$ for Mars in order to make sense of the meteorite CRE data; I will show that Warren's conclusions are justified by my simulations.

Substituting this form of Φ into Eq. 6.12 yields

$$N_l(p, E, v > v_e) = \int_{v_e}^{\infty} \mathcal{N}(v) dv = k \frac{D^{\alpha}}{p-1} \left(\frac{1}{v_e}\right)^{p-1} \; . \tag{6.16}$$

However, since at the moment we are interested only in a single body (the Moon), this simply says

$$N_l \propto D^{\alpha} \tag{6.17}$$

since all other quantities are fixed.

In light of the discussion of Sec. 3.4.1, let us quantify what we can say about the value of p (assuming a power law is valid) from the lunar meteorites. As discussed earlier, the majority of meteoroids must not leave the Moon at more than 3 km/s. It is easy to show from Eq. 6.16 that the fraction f_* of *escaping* meteoroids that leave at velocity *greater* than v_* is

$$f_* = \left(\frac{v_e}{v_*}\right)^{p-1} \; . \tag{6.18}$$

If at least 50% of the escaping lunar ejecta is to leave at less than v_*=3.0 km/s, then $p \geq 4$; if 90% then $p > 11$.

6.2.4 The cratering algorithm

I now proceed to simulate the cratering and delivery history from the Moon. The goal is to compile both the meteorite-time and impactor-size spectra over the entire impact history. That is, I will compute the expected CRE age spectrum (hereafter called the 'age spectrum') and the fraction of meteorites contributed by each size bin of the impactors (hereafter called the 'launcher spectrum', not to be confused with either the size distribution of the impactors or the size distribution of the ejecta). It should be noted that because there is a delivery 'window' of duration $A_{\rm ice}$, the algorithm below actually calculates the spectrum of *blast-off times* (Sec. 2.3.3), since all that is done is to compile the number of meteorites that have arrived in the last $A_{\rm ice}$ years.

The algorithm is as follows: for each size bin $(D, \sqrt{2}D)$, beginning with the largest and decreasing through the smallest $D_{\rm small}$, I perform the steps described below. The parameters are summarized in Table 6.1.

1. Compute the number of meteoroids launched by the impactor of this size, scaled from Eq. 6.17. I arbitrarily take 10^6 meteoroids to be launched by a 100-m impactor. This normalization is for convenience, its value makes no difference to the results.

2. For this bin, compute the number of impactors X_{imp} relative to the 1-km impactor bin, using Eq. 6.11. If desired, one can use a size distribution for $D < 100$ m with exponent $\beta \neq 2$, in which case one first computes the number of 100-m impactors and then scales from there using $D^{-\beta}$.

3. The *mean rate*, or average spacing in time, of impacts in this size bin D is thus $t_{space} = t_{peg}/X_{imp}$.

4. Now work backwards in time through the impact history. Each impact occurs randomly in time, and so the probability of *not* having had an event by time t is $\exp(-t/t_{space})$. Thus, generate a random number R from 0 to 1, and the next impact time is $t_{imp} = t_{space} * \ln(1 - R)$ farther back in time. When constraining the parameters I have sometimes turned off this random timing so that $t_{imp} = t_{space}$ always.

5. Use Eq. 6.3 to determine the number of meteoroids that this impact has delivered to the ice sheet during the recent past. Keep track of which size bin and time bin the particles are from, as well as the total number of lunar meteorites that have arrived. I have used the 2.4 km/s delivery spectrum to compute $\mathcal{F}(T)$.

6. Return to step 4, and continue generating events until $t > 100$ Myr. That is, assume that impacts occurring longer ago than 100 Myr are not delivering any appreciable number of meteorites today. (The lunar simulations in fact ended at 10 Myr, but I use the 100 Myr number to be consistent with the martian study to be done below.) I assume that only a negligible fraction of lunar meteorites return after 10 Myr; this is supported by both their absence in the collection and the relatively small delivery rate in the later stages of the simulations.

7. Once this impactor diameter class is finished, return to step 1 and repeat for impactors that are $\sqrt{2}$ smaller, until we reach the smallest impactor class D_{small} considered.

To summarize, the parameters relevant for the lunar simulation are given in Table 6.1, along with likely values for these parameters.

Table 6.1 Parameter values for studies of the lunar collision history. Reasonable ranges for each parameter are given; these values will be constrained later by comparison with the meteorite data.

Parameter	Description	Value
α	$\Phi \propto D^{\alpha}$	2 – 4
β	$N(D) \propto D^{-\beta}$ (for $<$ 100 m)	2 – 3
$A_{\rm ice}$	retention age of ice sheet	0.03 – 0.1 Myr
$D_{\rm small}$	smallest impactor bin	<100 m (maybe $\ll$?)
t_{peg}	Freq. D >1-km impactor hits Moon	4 Myr ($\pm$ factor of 3?)

6.2.5 Constraining the lunar parameters

Fig. 6.3 shows the baseline case; the top panel gives the 'age spectrum' and the bottom panel the 'launcher spectrum' (as defined above). The five relevant parameters are indicated (likely ranges for these parameters are given in Table 6.1). The notation $N(D) \propto D^{-2}, D^{-2}$ refers to the power laws for $D < 100$ m and $D > 100$ m, respectively; *i.e.,* the first number is β since the size exponent is always taken to be -2 above 100 m. The five parameter values given are in fact close to optimal for producing a good match with the data. I will vary the parameters away from this baseline case in order to demonstrate the various dependencies. For these purposes I temporarily will not allow the impactors to be randomly spaced, because this introduces variations in the spectra that obscure the systematic changes induced by altering the parameters. Random spacing will be reinstated in Sec. 6.2.5.5, but for now the impacts will occur spaced exactly by t_{space}; *i.e.,* if $t_{space} = 10^4$ yr, then the impacts are considered to occur 10^4, 2×10^4, 3×10^4 yrs, etc., ago. The properties of the size bins for the standard case are listed in Table 6.2.

Fig. 6.3 has several features of note. Two sets of meteorite data are shown, connected by dotted lines, and these agree above 250 kyr. The rightmost set of data uses only upper limits obtained from the available CRE data. The left set instead adopts the thermoluminescence (TL) 4π ages for those meteorites that have them (Table 2.1), and uses lower limits for all other quantities in order to calculate a minimum blast-off age. For Y-793274 and EET87521, which have no lower bound on t_{bo} available from either analysis, I have used a minimum t_{bo} of one-tenth the upper limit. It is likely that the left curve based on the TL data underestimates the actual transit ages; TL ages can be erroneously low due to heating of the meteoroid (see Benoit *et al.* 1996); the right curve serves as a definite upper bound. The correct spectrum is probably somewhere in the middle, like the computed spectrum in this case. I will continue to show both limits for the actual ages, and readers can judge for themselves.

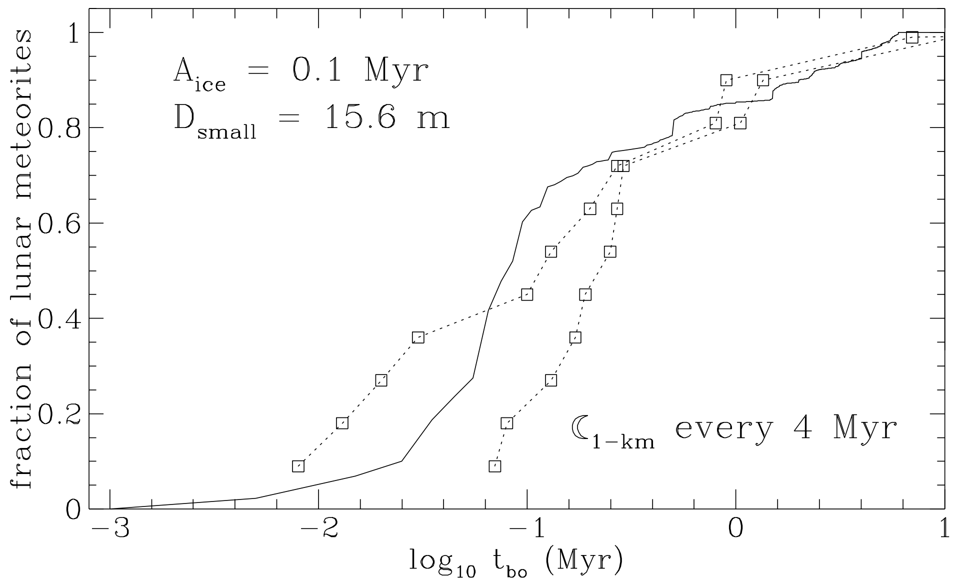

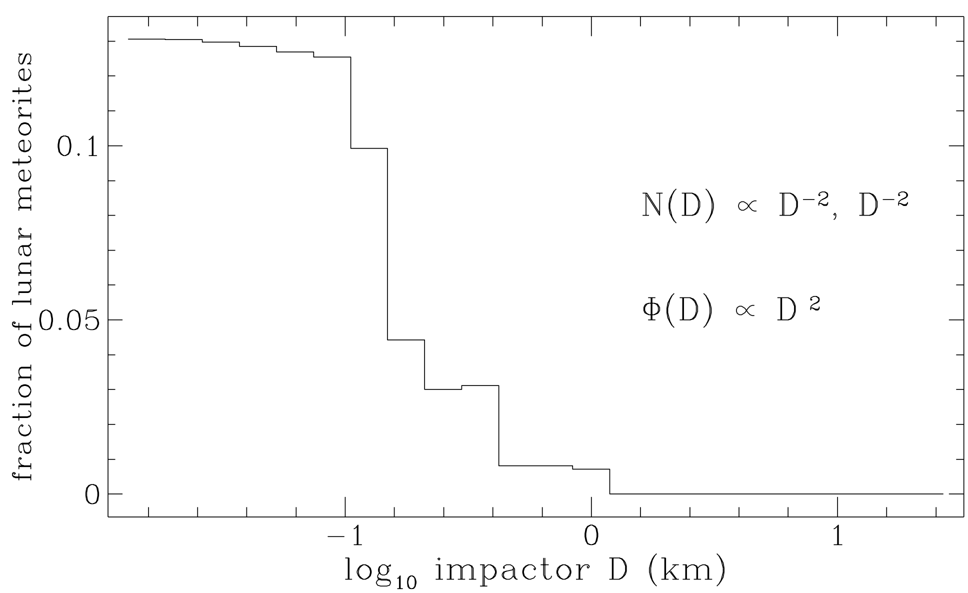

Figure 6.3 Upper panel: Age spectrum of simulated lunar deliveries (solid line) overlain with two sets of CRE data (points connected by dashed lines). The rightmost set of CRE data consists of upper limits from the cosmonuclide data; the left set is lower limits (see text). Lower panel: The launcher spectrum of the meteorites (the fraction of delivered meteorites launched by each size bin of impactors). In this case most of the meteorites in the ice sheet were launched by small ($D < 100$ m) impactors. The $N(D)$ notation indicates that a D^{-2} power law was used both above and below 100 m for the impactors. This calculation had the random timing of impactors turned off.

Table 6.2 Properties of the impactor bins for the lunar case. For each bin with lower bound $D_{\rm lower}$ (and upper bound $\sqrt{2}D_{\rm lower}$), the number of impactors X_{imp} relative to the $D_{\rm lower}$ =1-km bin is shown. The impactors of each bin hit the Moon on average t_{space} apart, and launch N_l meteoroids from the Moon above the escape velocity. The normalization for N_l was set at 10^6 meteoroids for a 100-m impactor, but the absolute value here has no effect on the spectra.

$D_{\rm lower}$ (m)	X_{imp}	t_{space} (Myr)	N_l
15.6	4.10×10^{03}	0.00098	2.44×10^{04}
22.1	2.05×10^{03}	0.00195	4.88×10^{04}
31.2	1.02×10^{03}	0.00391	9.77×10^{04}
44.2	5.12×10^{02}	0.00781	1.95×10^{05}
62.5	2.56×10^{02}	0.01563	3.91×10^{05}
88.4	1.28×10^{02}	0.03125	7.81×10^{05}
125.0	6.40×10^{01}	0.06250	1.56×10^{06}
176.8	3.20×10^{01}	0.12500	3.12×10^{06}
250.0	1.60×10^{01}	0.25000	6.25×10^{06}
353.6	8.00×10^{00}	0.50000	1.25×10^{07}
500.0	4.00×10^{00}	1.00000	2.50×10^{07}
707.1	2.00×10^{00}	2.00000	5.00×10^{07}
1000.0	1.000	4.00000	1.00×10^{08}
1414.2	5.00×10^{-01}	8.00000	2.00×10^{08}
2000.0	2.50×10^{-01}	16.00000	4.00×10^{08}
2828.4	1.25×10^{-01}	32.00000	8.00×10^{08}
4000.0	6.25×10^{-02}	64.00000	1.60×10^{09}
5656.9	3.12×10^{-02}	128.00000	3.20×10^{09}
8000.0	1.56×10^{-02}	256.00000	6.40×10^{09}
11313.7	7.81×10^{-03}	512.00000	1.28×10^{10}
16000.0	3.91×10^{-03}	1024.00000	2.56×10^{10}
22627.4	1.95×10^{-03}	2048.00000	5.12×10^{10}

A second interesting feature is that the launcher spectrum shows that even though large impacts that occurred in the past are included, almost all of the lunar meteorites now present on Earth were launched by impactors smaller than 1 km, and the majority by impactors <100 m. This lack of representation of large launchers is aided by, but not exclusively due to, α's value of 2 in this simulation (as discussed in Sec. 6.2.5.3 below) Notice a very important fact; even though more than half of the delivered meteorites from an impact arrive within 0.01 Myr of launch (Fig. 3.13) for the 2.4 km/s $\mathcal{F}(T)$ function we are using, the age spectrum in Fig. 6.3 predicts only $\simeq$5% of the simulated deliveries have blast-off ages this low. This rather obvious effect was not mentioned in Gladman *et al.* (1995); if the meteorites come from multiple impacts, only very recent (and therefore presumably smaller on average) impacts can deliver young samples to the ice sheet because meteorites that have landed in the last 50 kyr from impacts that occurred hundreds of kyr ago will not have short t_{bo} ages. This effect does not modify the argument (which is based solely on 4π ages) that the ejecta's velocity-spectrum must be steep (Sec. 6.2.3).

When judging the quality of matches to the data, we must keep in mind the small-number problem of the meteorite data. Accordingly, the most that should be required of the model is to reproduce the gross features of the problem to within a factor of 2 or 3. In the lunar case, these features are: (1) the age spectrum, with half of the meteorites arriving in less than 100,000 years; and (2) a lack of SCPs, indicated by having most of the meteorites launched by small impactors as opposed to a few large ones.

6.2.5.1 The smallest impactor $D_{\rm small}$

Fig. 6.3 showed the case where an impactor between 1 and $\sqrt{2}$ km hits the Moon every 4 Myr. Since the impactor size distribution varied as D^{-2} below 100 m, impactors between D=88 and 125 m hit the Moon roughly every 30 kyr (Table 6.2). As Warren (1994) points out, impactors capable of launching lunar meteorites must be smaller than this since we have had several launches in the last 50 kyr (barring a statistically unlikely barrage of recent impactors of this size range). This raises the question: how small an impactor can launch a D=2 cm, $\simeq$10 g lunar meteorite? This is a difficult question to answer beyond the trivial argument that the impactor would have to be at least of order the same diameter. Can we perhaps discover the answer from the meteorite record?

Fig. 6.4 shows how the expected meteorite 4π-age spectrum would change as a function of $D_{\rm small}$. The predictions are essentially identical for $D_{\rm small}$ <31 m, but above this the spectra begin to not agree with the data. This is because below $D_{\rm small}$ =63 m the smallest impactors, which deliver most of the meteorites, strike much more often than $A_{\rm ice}$ (see Table 6.2). Above $D_{\rm small} \simeq$ 63 m few, if any, impactors strike the Moon more recently than $A_{\rm ice}$, and therefore almost no meteorites with t_{bo} < 0.05 Myr are present. The (62.5,88.4) m impactor bin

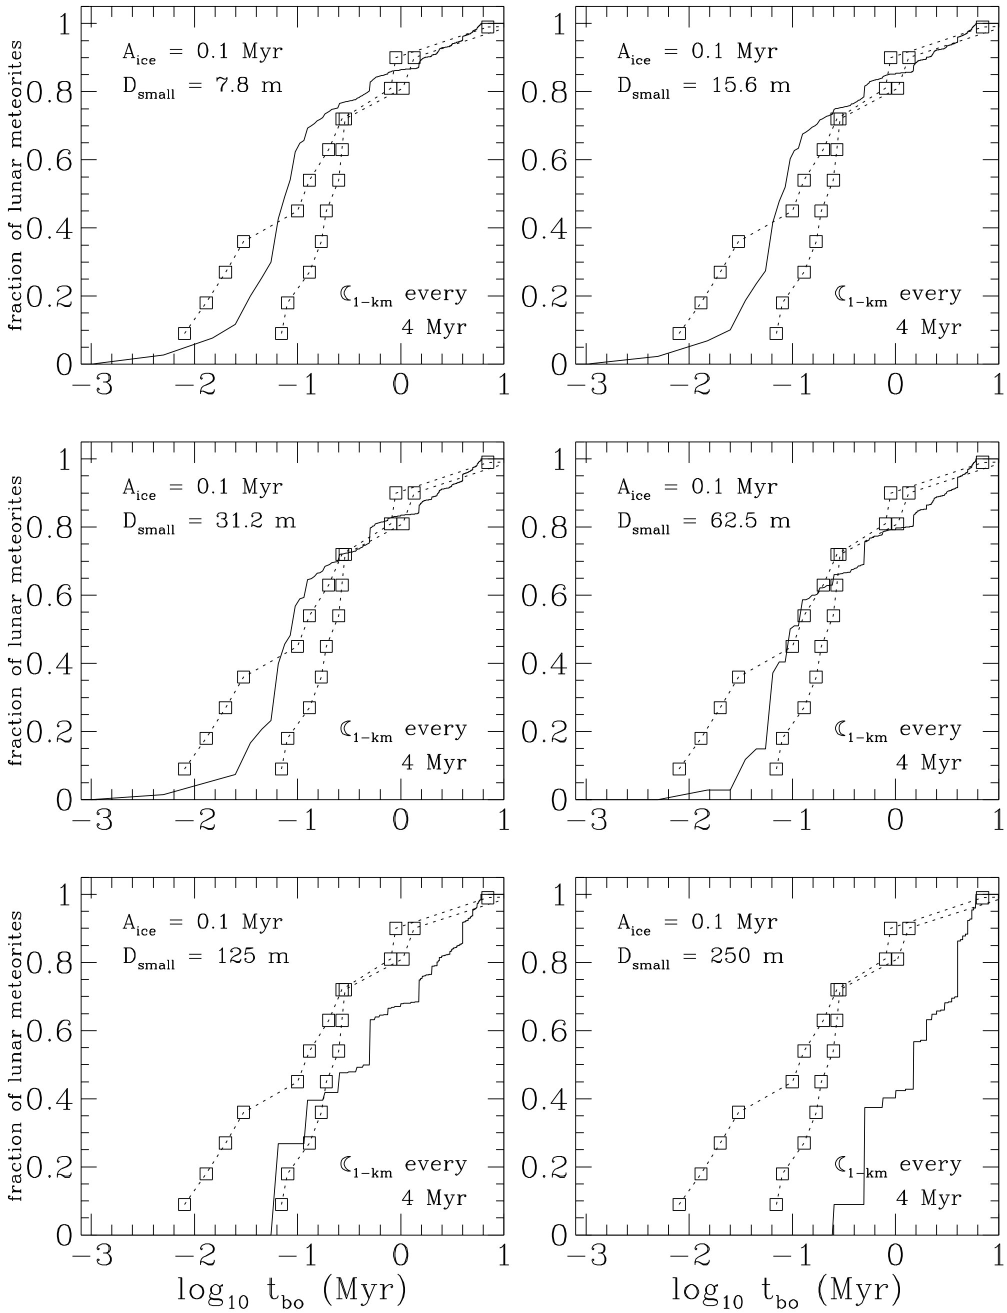

Figure 6.4 These panels show the expected age spectra as a function of the smallest diameter impactor that is allowed to launch lunar meteorites. The launcher spectra (not shown) each look similar to that of Fig. 6.3, in that the majority of the delivered meteorites are launched by impactors smaller than $D \simeq 300$ m.

delivers an impactor to the Moon every 16,000 years, so apparently the impactors capable of launching meteoroids must be striking at least this often in order to produce the observed age spectrum. Unfortunately, the age spectrum is not sensitive to a minimum value for D_{small}, so this remains unconstrained. However, below 63 m the value of D_{small} does not effect the results.

6.2.5.2 The age of the ice sheet A_{ice}

Varying the assumed average retention age of the ice sheet for meteorites results in the expected spectra changing as shown in Figs. 6.5 and 6.6. As A_{ice} increases, the fraction of young meteorites drops monotonically. The age spectrum can be seen to be fairly sensitive to the value of A_{ice}; an older ice sheet becomes more able to retain meteorites from larger, more ancient impacts. In fact, only in the range A_{ice}=0.08–0.16 Myr does the predicted age spectrum agree with the data (by 0.32 Myr the spectrum is too close to the upper limits of the CRE data to be considered a likely match). This age range for A_{ice} is entirely consistent with the discussion of 6.1.4.1, and so from now on I shall adopt A_{ice}=100 kyr, in agreement with the lifetime of Antarctic meteorites determined by radioactive dating (Sec. 6.1.4.1).

Fig. 6.6 is interesting because it clearly shows the effects of an unrealistically long retention age for the ice sheet. Even though it is still the case that almost all of the meteorites are launched by small impactors, only a tiny fraction of those meteorites will have blast-off ages $<$0.25 Myr due to the fact that there are so many more small impacts between 0.25 and 1.0 Myr that have delivered their meteorites to the long-lived ice sheet. Thus, the preponderance of short blast-off times in the lunar meteorite sample is due to two effects: the very high initial accretion rates due to low ejection velocities *and* the young ice sheet.

6.2.5.3 The launch exponent α

The value of α has a very large effect on the incoming meteorite flux. If $\Phi(D) \propto D^2$ (*i.e.,* α=2) the number of meteoroids launched per impact scales as the surface area of the crater since the crater diameter is roughly proportional to that of the impactor (Melosh 1989). This would be expected if the meteorites are being launched in a very thin surface layer of roughly constant thickness, as both the CRE evidence and the theoretical considerations of the spallation model would predict. On the other hand, $\alpha = 3$ might be expected if the number of meteoroids launched scales as the volume of the crater (Warren 1994), with a spall zone thickness $\propto D$. Fig. 6.7 shows that if we adopt this latter scaling, the arriving meteorite flux would be inconsistent with the data.

An important property of the $\Phi(D) \propto D^2$ scaling is that, if the impactor distribution scales as $N(D) \propto D^{-2}$, then each logarithmic bin will *launch* equal numbers of meteoroids. The arriving spectrum will thus be entirely controlled

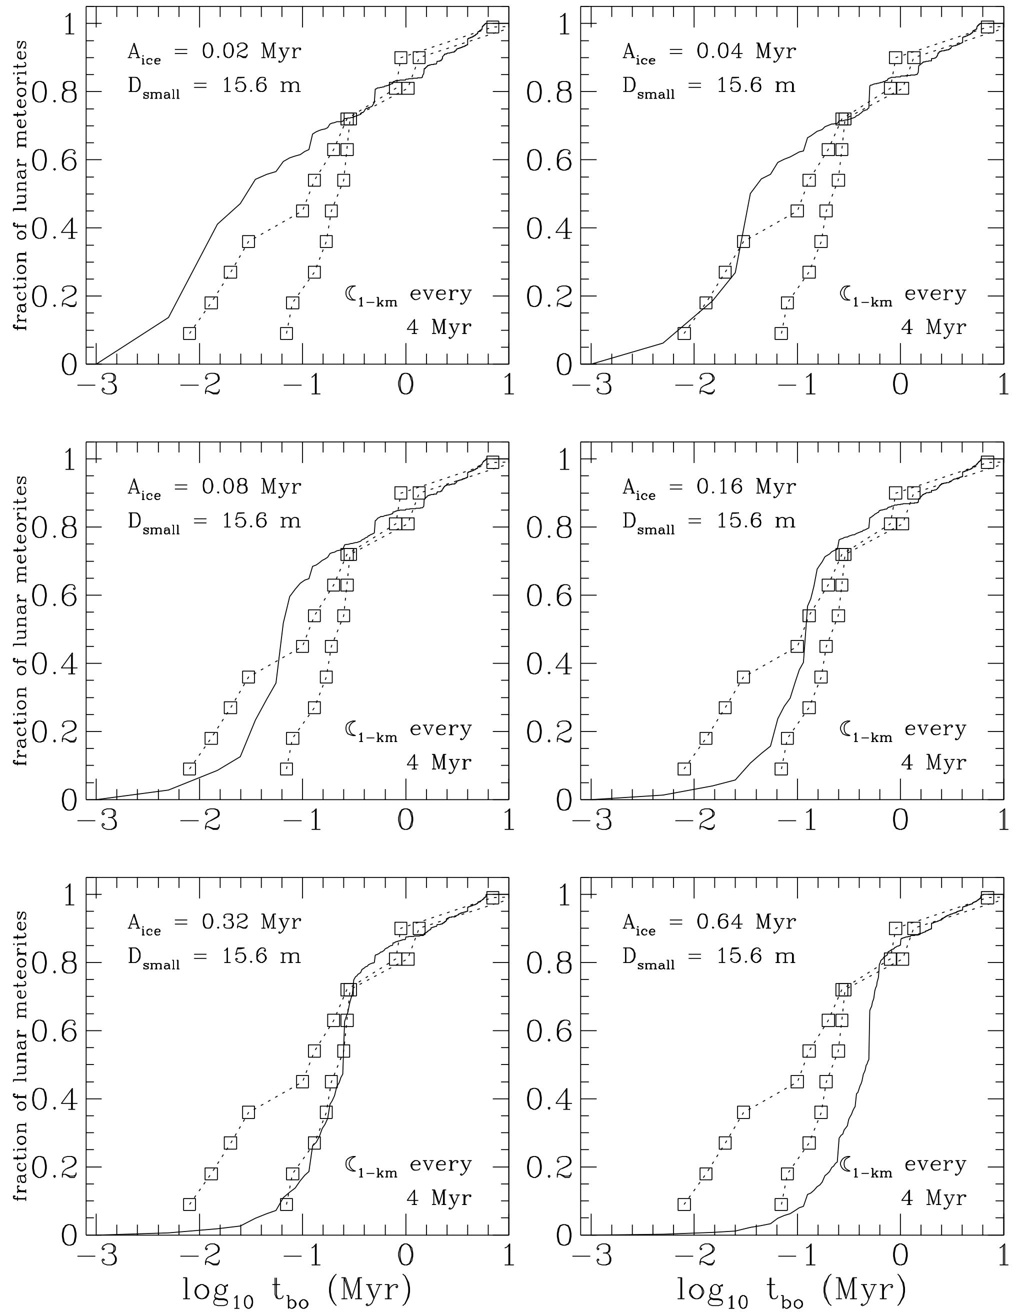

Figure 6.5 These panels show the expected age spectra as a function of the retention age of the ice sheet for lunar meteorites. The launcher spectra (not shown) again look similar to that of Fig. 6.3, in that the majority of the delivered meteorites are launched by impactors smaller than $D \simeq 100$ m.

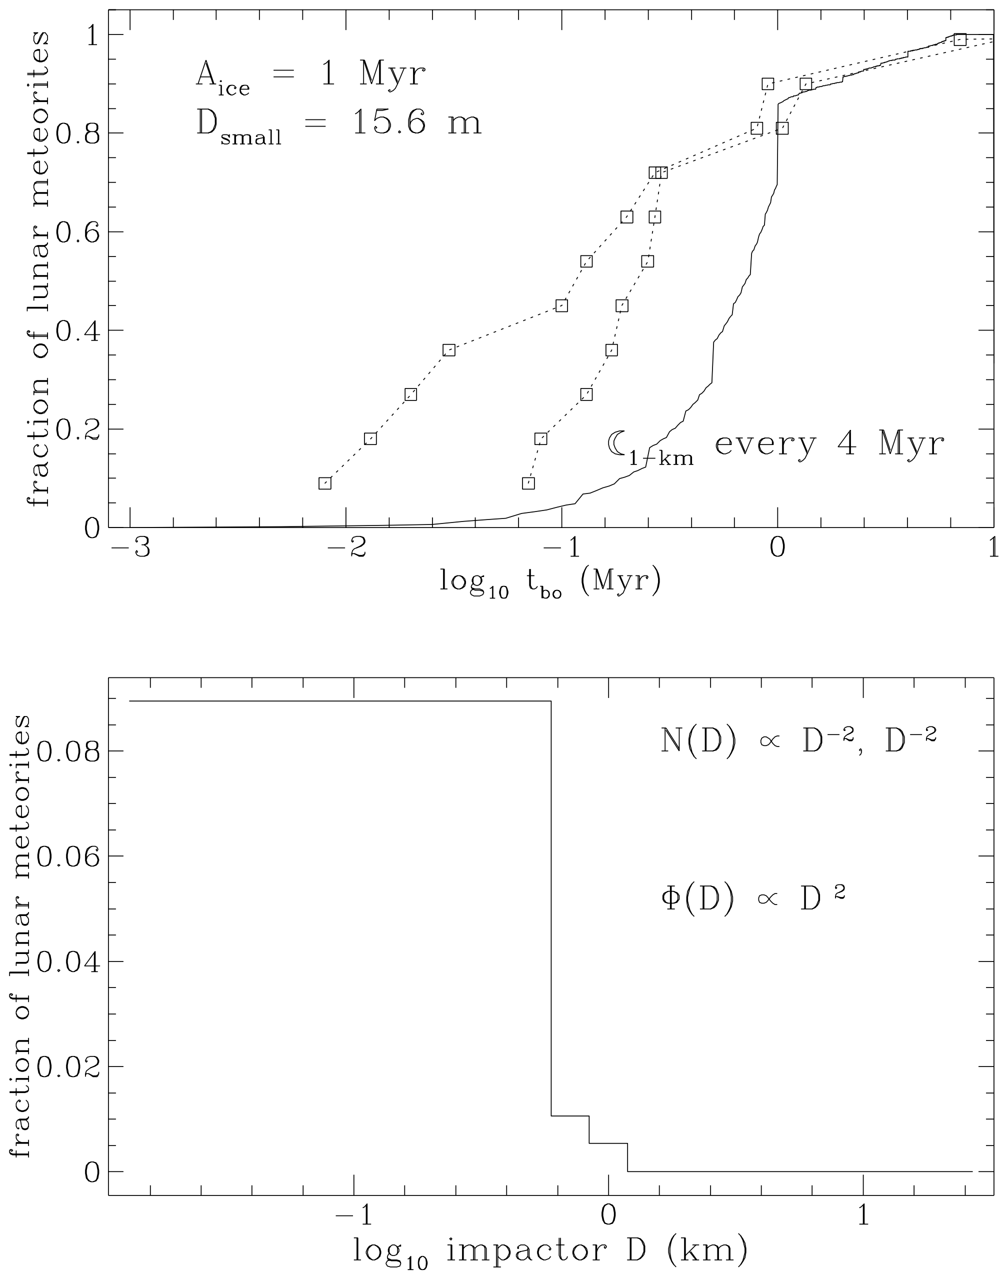

Figure 6.6 The age and launcher spectrum for the case of A_{ice}=1 Myr, corresponding to an extremely old ice sheet. Using such an unrealistically long retention age results in almost all of the meteorites presently in the ice sheet coming from impacts that occurred more than 0.25 Myr ago; this is firmly ruled out by the CRE data.

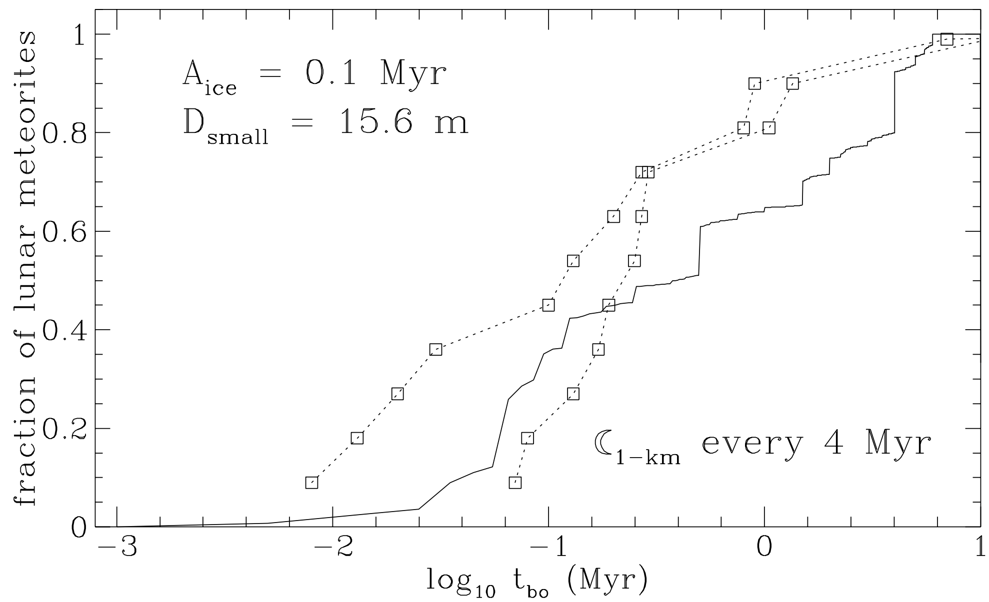

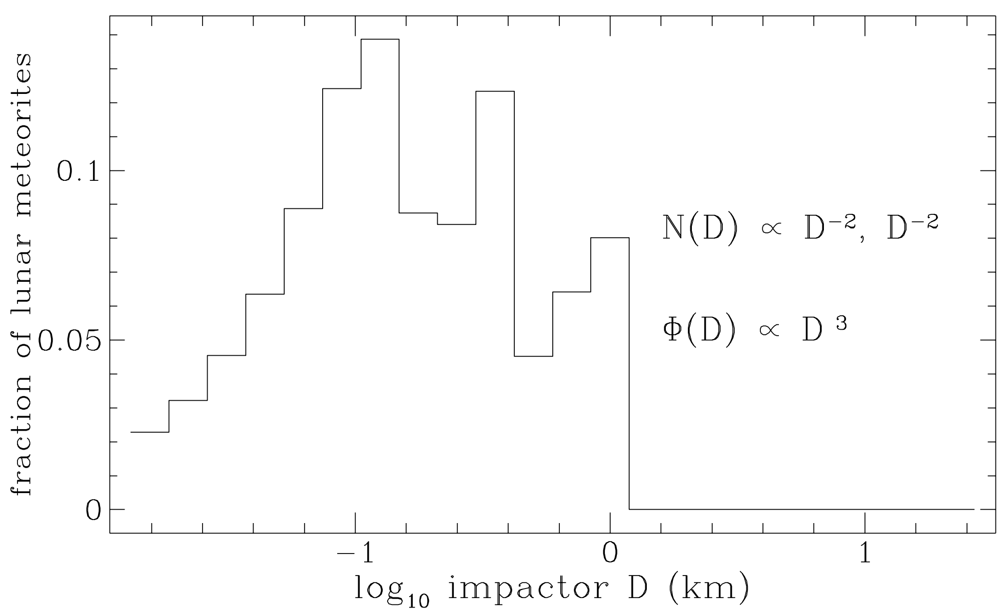

Figure 6.7 The age and launcher spectrum if the number of meteoroids launched per impact scales as the cube of the impactor diameter. Too many of the meteorites should then come from more than several hundred kyr ago, in disagreement with the data.

by the $\mathcal{F}(T)$ spectrum and the influence of the ice sheet. The drop-off in the launcher spectrum beginning at $D \sim$100 m in the standard case of Fig. 6.3 is due to the fact that impactors larger than this are not hitting the Moon on a time scale less than $A_{\rm ice}$. The absolute cut-off at $D \simeq$1 km in that same figure results from no impactors larger than 1 km having hit the Moon in the $\sim$10 Myr time scale over which lunar meteorite delivery takes place (*i.e.,* where $\mathcal{F}$ becomes close to 1). This reasoning is further clarified by a re-examination of Fig. 6.6, where all impactors that hit more often than the age of the ice sheet are equally represented, and the very dramatic cutoff still occurs at $D \simeq$1 km.

When α increases to 3, this behavior changes. Now the larger impactor bins launch more meteorites than the smaller ones, but these larger impactors occur less frequently, and thus on average longer ago. These two effects produce a launch spectrum as shown in Fig. 6.7, with most of the meteorites coming from impacts of objects a few hundred meters in diameter. The cut-off at 1 km is of course unaffected. If α were in fact 3, then not only would the age spectrum of meteorites be different, but we should probably expect a greater prevalence of more deeply launched meteorites since larger impactors should excavate to greater depths. It is notable that the three oldest lunar meteorites are also the ones that seem to have been launched from the greatest depths (Table 2.1, p. 16).

6.2.5.4 Steepening the size distribution $N(D)$

As long as α=2, steepening $N(D)$ below 100 m to be proportional to D^{-3} leaves the age spectrum virtually unaffected. This occurs because the majority of meteorites already come from small impactors spread evenly in time; thus adding additional smaller impactors occurring more often than $A_{\rm ice}$ only slightly modifies the age spectrum, producing a small enhancement in early deliveries (Fig. 6.8). The launcher spectrum is dramatically affected, however, since the extremely numerous impactors in the smallest size bins launch most of the meteorites. Since we have no direct evidence of what size impactors launched the lunar meteorites (although shallow launch depths probably imply smaller impactors), the lunar meteorite data cannot constrain the possibility that the impactor size distribution steepens for smaller bodies. However, neither do we have to worry about the possible steepening affecting the model either, for the age spectrum is insensitive to this change.

The absolute number of meteorites will depend rather strongly on this steepening. An $N(D) \propto D^{-2}$, $\Phi(D) \propto D^2$ distribution has the property that adding another size bin at the lower end increases the total number of launched meteorites only by a small amount since every bin contributes the same number of meteorites. On the other hand, if the size distribution steepens to D^{-3}, then each additional bin contributes more meteorites than the next larger one, explaining the skewing of the launcher spectrum towards smaller impactors.

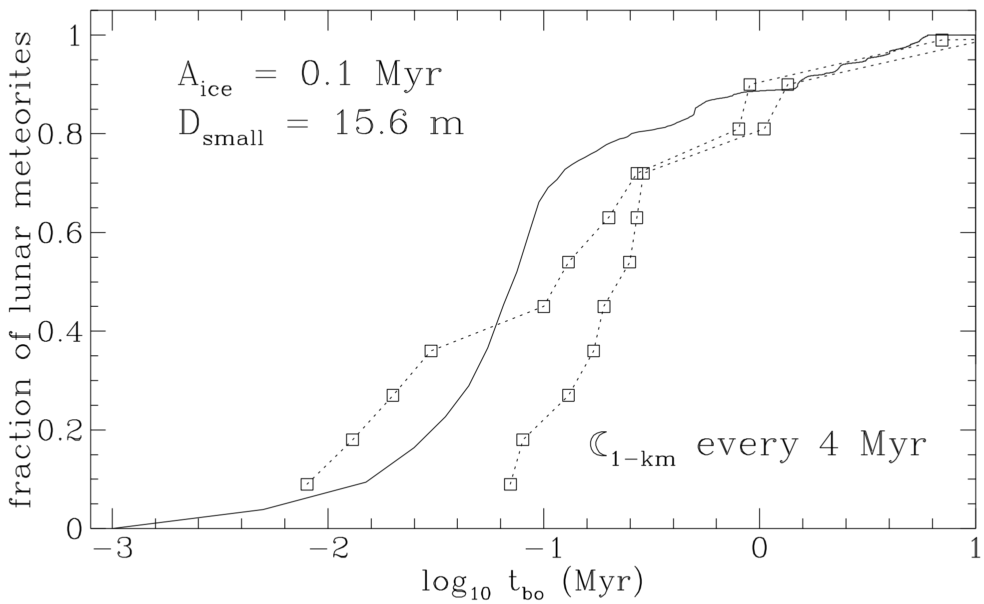

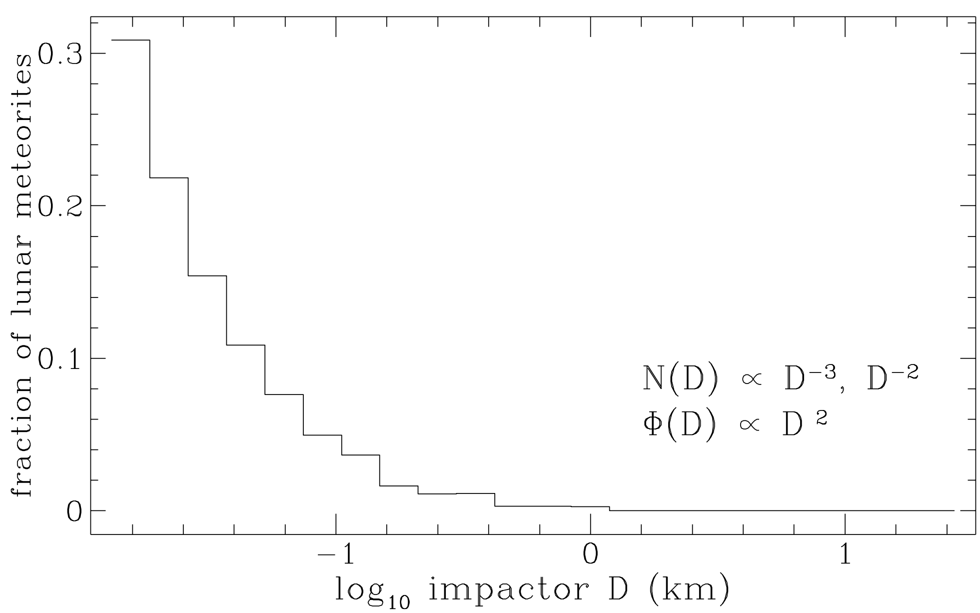

Figure 6.8 The age and launcher spectrum if the size distribution of the impactors steepens to $N(D) \propto D^{-3}$ below 100 m. The age distribution is skewed very slightly towards faster returns (compare to Fig. 6.3). The launcher spectrum changes more dramatically (see text).

6.2.5.5 Lunar source-crater pairings

One obvious way to produce an excess of younger meteorites (the chief constraint we have used to match the age spectrum) is to have a recent large impact. To illustrate this, now that we have explored the systematic effects of changing the other parameters, I have included the random spacing of the impactors (step 5 of the algorithm in Sec. 6.2.4), which allows larger impactors to have struck the Moon in the recent past. Fig. 6.9 shows such a case; a single 800-m impactor struck the Moon $\simeq$ 40,000 years ago. As can be seen from the launcher spectrum, this impact launches so many meteorites, and has such a high transfer efficiency, that about half of all the meteorites expected to be in the ice sheet are contributed by this single impact. However, all of these meteorites should have identical blast-off ages and similar petrologic characteristics (*i.e.,* they should be source-crater paired). Since Warren (1994) feels that few, if any, of the lunar meteorites are source-crater paired, we can rule out the occurrence of any large ($\gtrsim$300 m) impacts on the Moon in the last hundred thousand years. Note that the effect of these impactors would be more pronounced if we were to allow α=3, since the larger impactors would launch even more meteorites.

This example highlights the importance of the stochastic nature of the impact process. By performing about 30 such simulations I have observed that an impactor large enough to produce an unacceptable number of source-crater pairs occurs only in about one-third of the cases. Even if an impact larger than 500 m happened in the last 10 Myr, it can occur long enough ago that its initial peak of deliveries is prevented from 'contaminating' the age spectrum as illustrated by Fig. 6.10. I suggest that the source-crater pairing of Asuka-881757 and Y-793169 may be due to the occurrence of such a large impact $\simeq$1 Myr ago on the Moon; the deep provenance of these two meteorites supports such a hypothesis.

6.2.6 Lunar summary

In summary, of the five relevant parameters, only $A_{\rm ice}$ and the exponent α of $\Phi(D)$ dramatically affect the spectra (for plausible ranges of the parameters). Apparently α must be closer to 2 than 3, and thus the number of meteoroids launched per impact scales with the surface area of the crater rather than its volume; this conclusion was already reached by Warren (1994) on the basis of the number of recent launch events. The modelling suggests that the ice sheet is less than about 60 kyr old, which is perhaps somewhat younger than one might estimate (see Nishiizumi *et al.* 1989); however, recall that the relevant parameter $A_{\rm ice}$ is actually the mean survival time of lunar meteorites on the ice surfaces that have been searched so far. If we base our estimate of $A_{\rm ice}$ on the meteorite sample (Sec. 6.1.4.1), then 60 kyr is reasonable.

The diameter of the smallest impactors capable of launching lunar meteoroids, the mean impact rate of 1-km impactors, and a possible steepening of the size

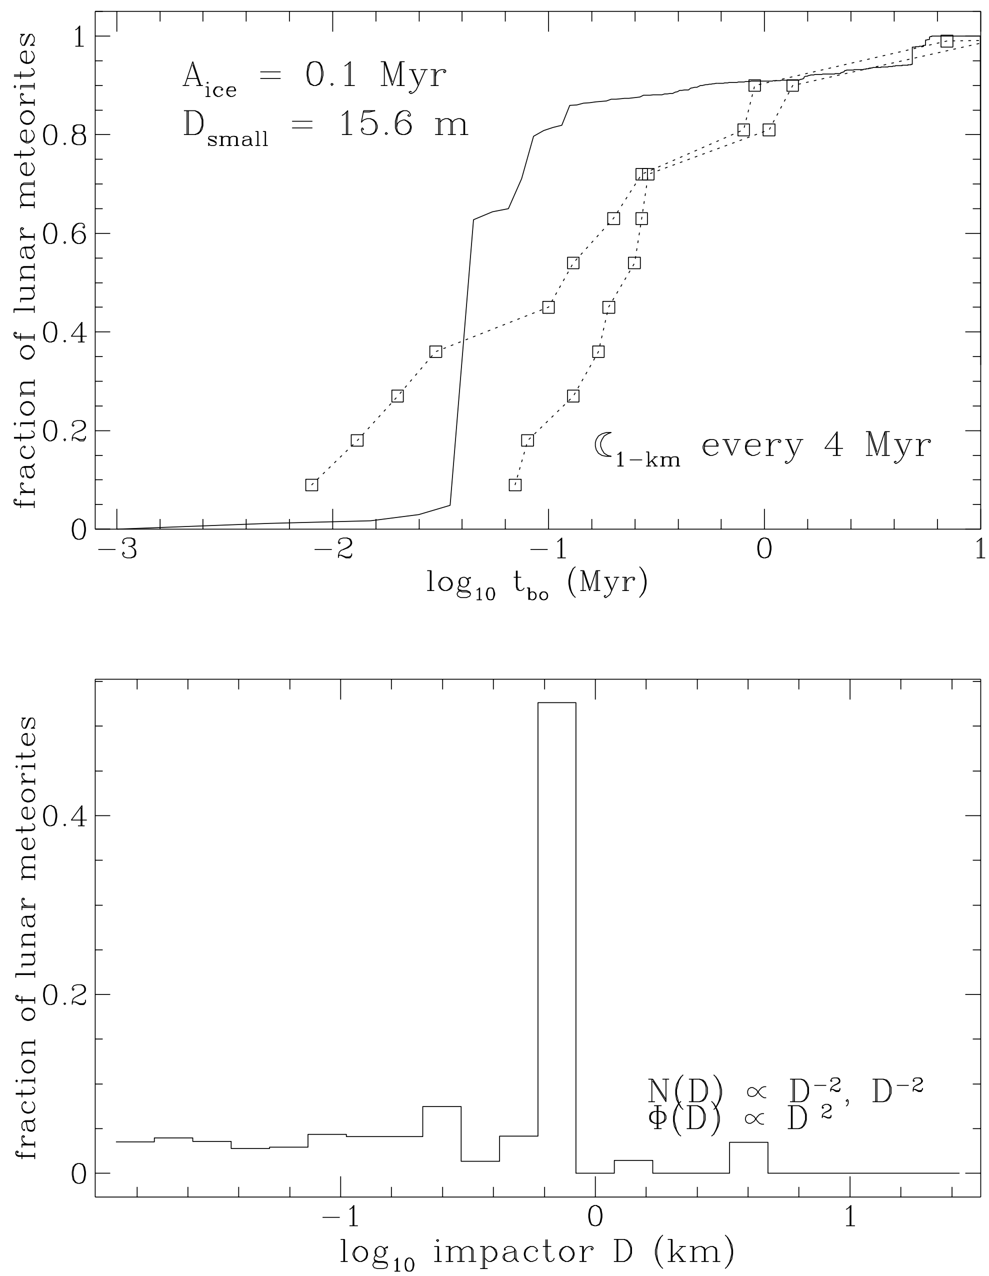

Figure 6.9 With random impact timing allowed, it is possible to have a large impactor hit the Moon recently. In the case shown, a single 800-m impactor striking the Moon 40,000 years ago would contribute about half of the lunar meteorites present in the ice sheet. The impact is indicated by the large spike in the launcher spectrum and the discontinuous jump in the age spectrum. This possibility is not supported by the meteorite data.

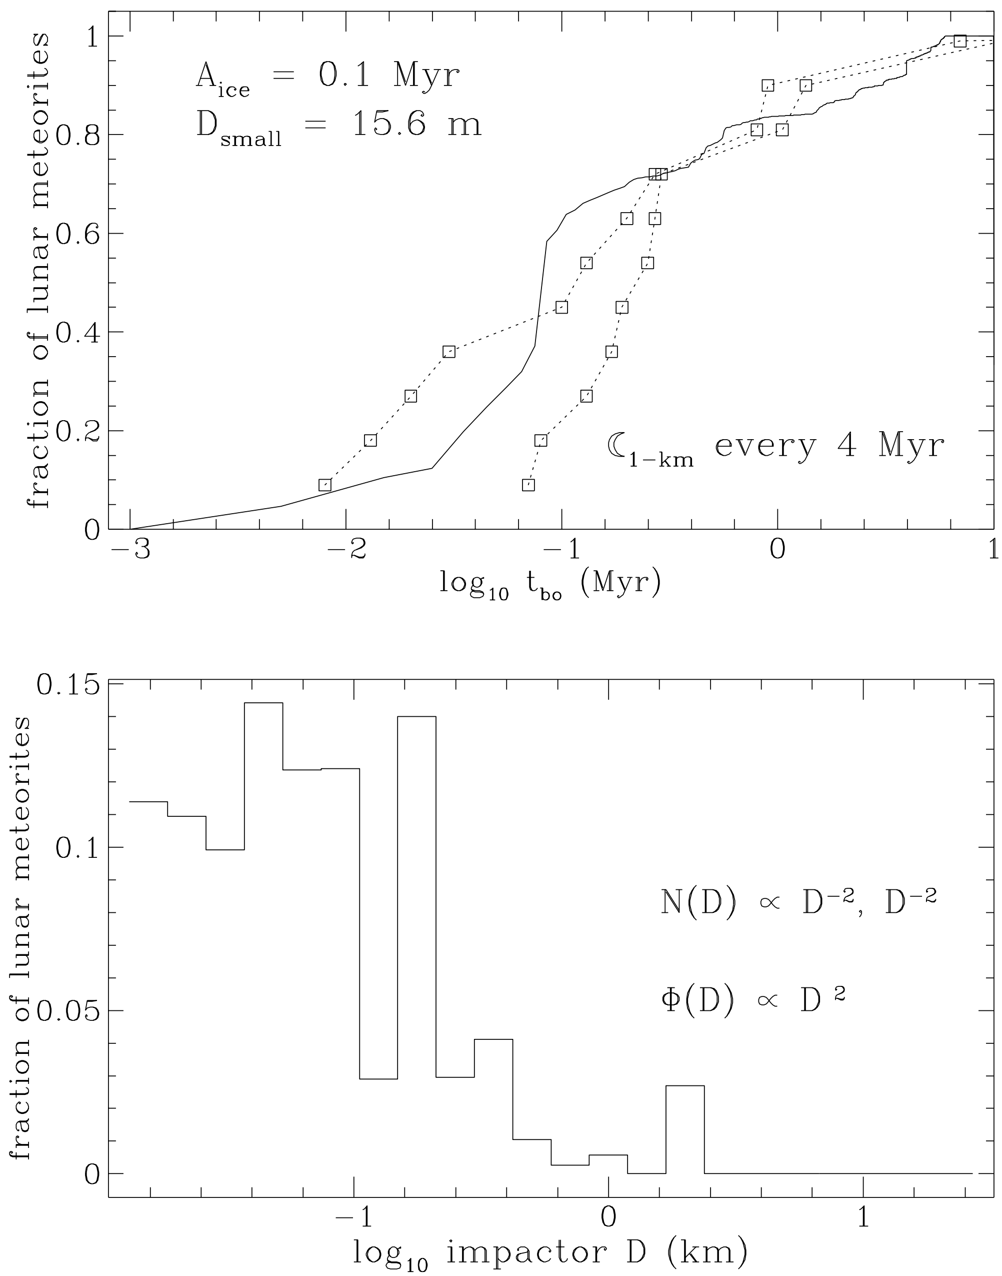

Figure 6.10 If a large impactor strikes the Moon more than several hundred thousand years ago, it delivers so little material to the ice sheet in the recent past that it produces only a minor effect in the age spectrum. In the case shown a 2 km impactor strikes the Moon $10^{0.6} = 4$ Myr ago (its presence is indicated by the small vertical jump of the age spectrum at this time). The peaks in the launcher histogram at large D are due to individual impacts, whereas at small D the meteorites are contributed by a myriad of small impacts.

distribution of the impactors below 100 m all have relatively small effects on the age spectrum.

6.3 Modelling martian launch

Our chief concern once we broaden our study to include the launch of martian material is to understand the circumstances under which the number of lunar to martian meteorites is roughly equal. This requires the inclusion of four additional parameters: (1) the relative flux of impactors hitting Mars as compared to the Moon; (2) the exponent p in the scaling of the number of launched meteoroids vs. ejection velocity in Eq. 6.12; (3) the energy exponent of $\Phi(D)$ for Mars; and (4) the fraction of the martian surface suitable for meteoroid ejection.

The flux of impactors onto Mars is not well constrained observationally. Theoretical models estimate that the space density of impactors near Mars should be comparable or slightly larger than that near the Moon. Wetherill (1989) calculates that the cratering rate *per unit area* on Mars is 1.94 times that of Earth. One might naively expect that the space density of impactors should be higher at Mars since it is nearer to the main asteroid belt, but the population of impactors in near-Mars and near-Earth space depends critically on the processes which inject asteroids and comets (the two classes of impactors) onto orbits inside the asteroid belt. However, there is no particular reason to assume that the *size* distribution of impactors in these classes should be any different at the Moon and Mars. Thus I assume that either asteroids or comets dominate at both targets, and that the flux of impactors $I_{\rm Mars}(D)$ near Mars is some fixed multiple of the lunar flux irrespective of size; that is,

$$I_{\rm Mars}(D) = F I_{\rm Moon}(D) \propto F D^{-2} \ . \tag{6.19}$$

If the space density of impactors were identical and gravitational focussing were ignored, then F would be simply the ratio of the martian to lunar surface area, which is four. The gravitational focussing of Mars is $\propto 1 + (v_{esc}/v_\infty)^2$ (Eq. 3.4). Most Earth-Moon impactors approach with v_∞ = 10–20 km/s (Wetherill 1969, Chyba *et al.* 1994, *cf.* p. 94) and, since we expect the martian values to be comparable, gravitational focussing likely results in an enhancement of a few tens of percent in the martian impactor flux.

I use F = 8, which combines focussing, Wetherill's estimate, and an order unity enhancement in the space density of impactors; F is likely to be accurate to a factor of two.

The assumption of a D^{-2} power-law over the entire impactor size range is probably even more valid for Mars since one would expect its higher escape velocity (of 5.03 km/s, more than twice that of the Moon) to result in a larger minimum impactor size for meteoroid launch. From Eq. 6.16, the minimum diameter for launching a meteoroid can be seen to scale as $v_{esc}^{(p-1)/\alpha}$, where p is the exponent of the velocity spectrum of the escaping particles. Since we concluded earlier that $p \geq 4$ (Sec. 6.2.3), martian launch requires impactors at least 3 times larger (and thus an order of magnitude less frequent) than lunar ones, assuming similar surface properties ($\alpha = 2$).

The exponent p is also very important for determining the relative number of launched meteorites. If surface properties were identical, Eq. 6.13 predicts that for *equal size impactors*,

$$\left.\frac{N_{Moon}}{N_{Mars}}\right|_{\text{launched}} = \left(\frac{v_{Mars}}{v_{Moon}}\right)^{p-1} \simeq 2^{p-1} \tag{6.20}$$

in terms of the ratio v_{Mars}/v_{Moon} of the escape speeds. The number of meteorites that arrive at Earth is influenced by the delivery efficiencies. If we adopt 40% for the lunar and 4% for the martian total delivery efficiency (Sec. 6.1.4), then the ratio of numbers that eventually reach the Earth from impactors of equal size is

$$\left.\frac{N_{Moon}}{N_{Mars}}\right|_{\oplus} = 10 \times \left(\frac{v_{Mars}}{v_{Moon}}\right)^{p-1} \sim 2^{p+2}\,. \tag{6.21}$$

Even if p is the probably minimal value of 4 from the lunar data, then the number ratio should be at least 60, as opposed to the observed value of $\simeq$1.

I will allow the exponent α of $\Phi(D)$ for Mars to be different from that of the Moon. This was done to permit the possibility of a general difference of surface properties between Mars and the Moon. Warren (1994) claims that the martian α must be closer to 3 than 2 in order to match the CRE evidence; I will explore the consequences of this. Lastly, I will also examine the effect of limiting the available area on Mars for meteoroid launch, in order to help reproduce the prevalence of young source terrains for the martian meteorites.

6.3.1 The baseline martian case

Fig. 6.11 shows the baseline case for the full simulation that considers both Mars and the Moon. The match here with the meteorite data is very good; large numbers of source-crater pairs are expected for the martian meteorites, and the ratio of lunar to martian meteorites R (indicated in the bottom left plot) is within a factor of 2 of unity. The exact timing of the large martian impact events need not precisely agree with the data of course; only the general spectral shape and prevalence of source-crater pairing need to be reproduced. The values of p, F, and the exponent of $\Phi(D)$ are also indicated on the martian panels. I shall again vary the parameters systematically to explore the model dependencies.

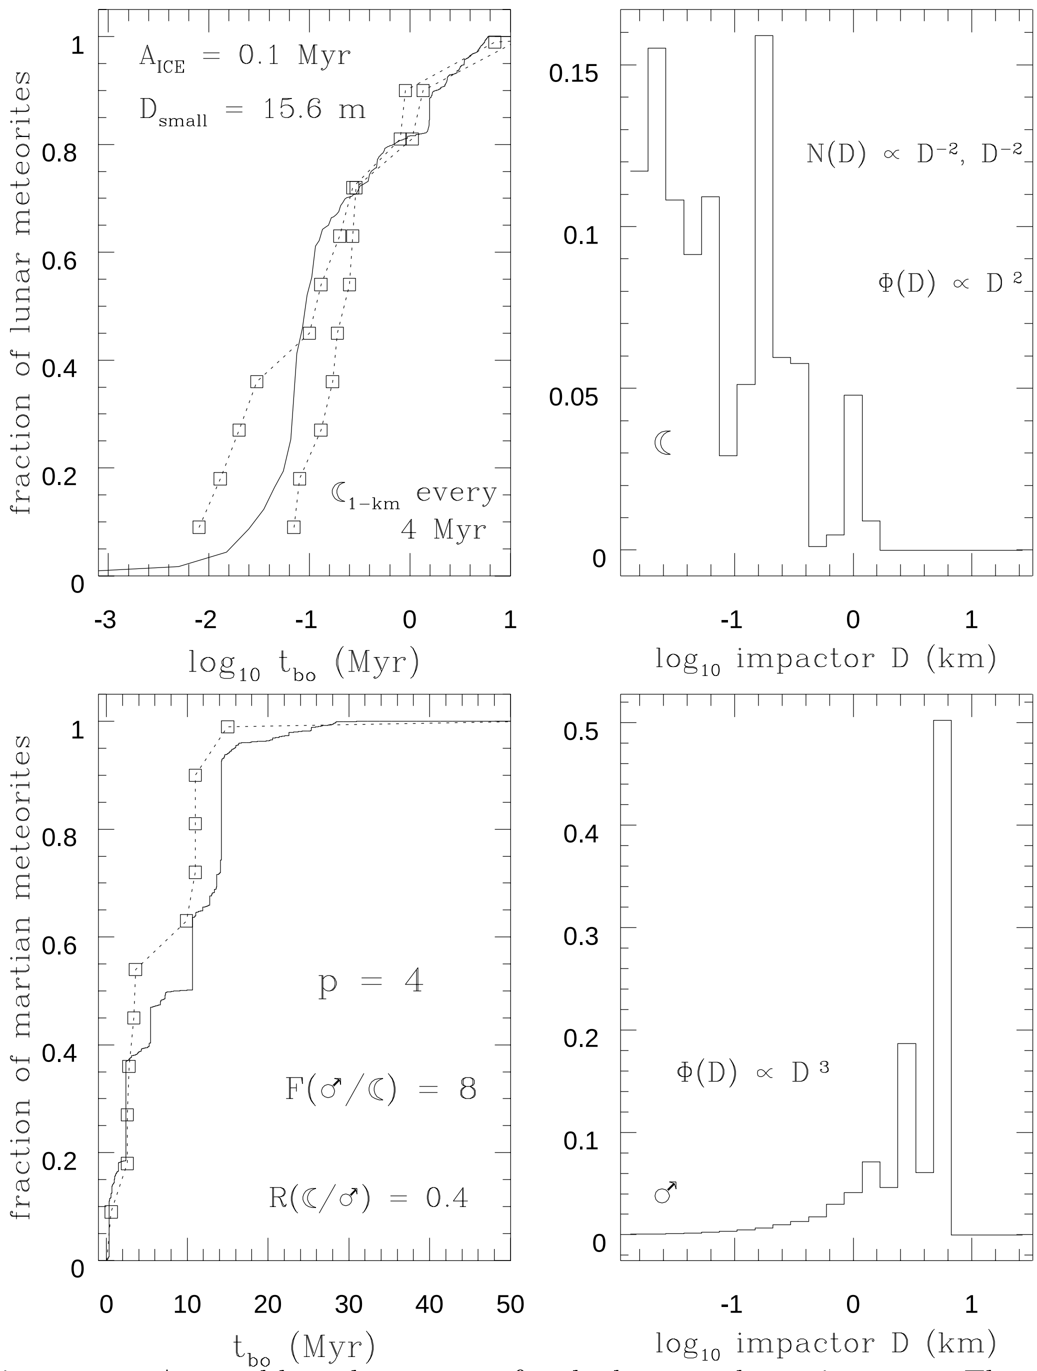

Figure 6.11 Age and launcher spectra for the lunar and martian cases. The two left panels show the age spectrum for the lunar (top) and martian (bottom) meteorites; note that the lunar panel uses a logarithmic time axis whereas the martian one is linear. The two right panels show the lunar (top) and martian (bottom) launcher spectra. *Random impact timing was included in this simulation.*

6.3.2 The velocity power law and impactor flux ratio

The only way in which the exponent p of the velocity power-law (Eq. 6.12) enters into the model calculation is when determining the ratio of the numbers of meteoroids launched in an impact of a given diameter (see Eq. 6.20). Thus, when p is increased by one, the lunar to martian meteorite ratio R goes up by a factor of about two (the ratio of the escape velocities). Thus changing p to 7 would not alter Fig. 6.11 at all except to change the meteorite number ratio to $\simeq 8 \times 0.4 = 3.2$.

Changes in the value of F produce an approximately linear change in the value of R since the flux of martian impactors in each bin is proportional to F. The model is basically incapable of distinguishing between variations in the parameters F and p, but I have kept both since they represent different physics.

6.3.3 The martian launch exponent

The martian spectrum depends critically on the value of α. In order to explore the effect of α's variation, I again temporarily turn-off the random spacing of the impactors (Fig. 6.12). As discussed in the context of the lunar case, if α=3 the larger impactor bins launch more material than the smaller ones, as shown in the martian launcher spectrum. Thus the largest impactor ($D \simeq 5$ km), which struck Mars $\simeq$16 Myr ago in this simulation, contributes the most meteorites (slightly more than 30%). All of the (source-crater paired) meteorites from this event would have the same blast-off age, as indicated by the discontinuity in the martian age spectrum. In this simulation the ratio R of lunar to martian meteorites is 0.6. A factor of two more martian meteorites have fallen during the age of the ice sheet, even though more lunar than martian meteorites have been delivered to Earth over all of history. This occurs because the current delivery rate from ancient lunar impacts is reduced to a relative trickle, whereas any martian impact delivers meteorites at almost a constant (but smaller) rate for $\approx$15–20 Myr (Fig.6.1). Since large impacts on Mars are more likely (due to the larger impactor flux), ancient martian impacts that launch large numbers of meteoroids are heavily represented in the meteorite record; this is doubly important because only when sampling these large impacts is there a reasonable chance of seeing source-crater pairs.

When α for Mars is changed to 2 (Fig. 6.13) the martian launcher spectrum is changed radically. Here, as in the lunar case, almost all the meteorites come from small impactors, because α=2 implies that each logarithmic bin launches the same number of meteoroids. Note from the expected martian age spectrum that although the distribution of transfer times is roughly correct, source-crater pairings are almost absent in the expected age spectrum. This disagrees with the meteorite data. Observe also that the lunar to martian meteorite ratio has risen by an order of magnitude.

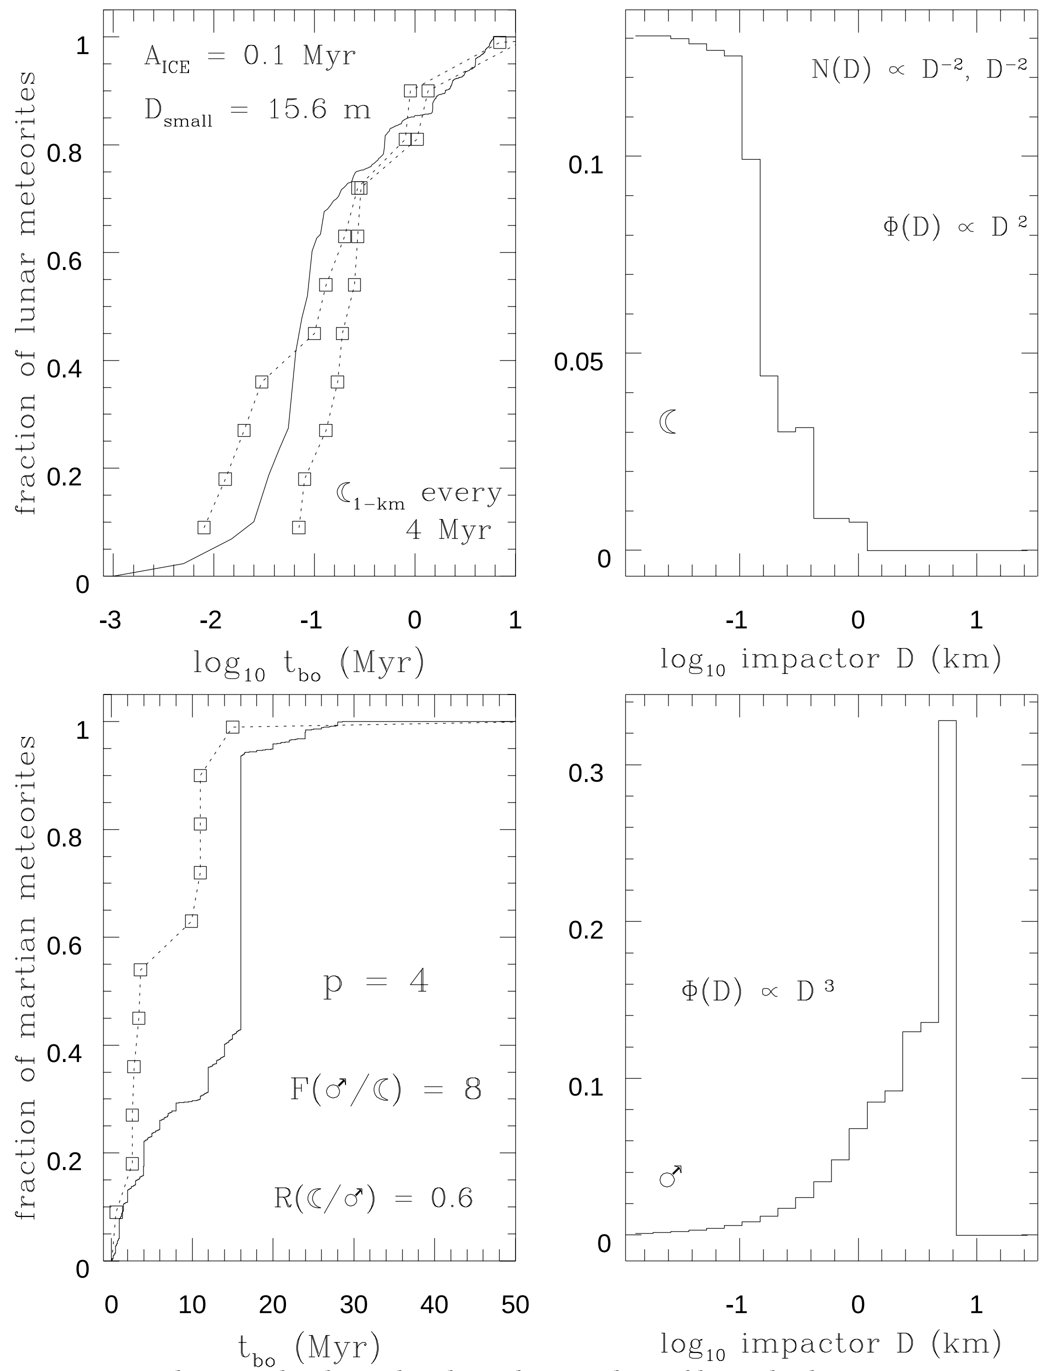

Figure 6.12 The standard result when the number of launched martian meteorites scales as the volume of the crater and random impacts are *not* allowed. The age spectrum of the martian materials is unrepresentatively skewed towards later launches. Note that the D^3 scaling of the launchers means that most of the martian meteorites are launched by the largest impactors that occur during the $\approx$20 Myr delivery horizon.

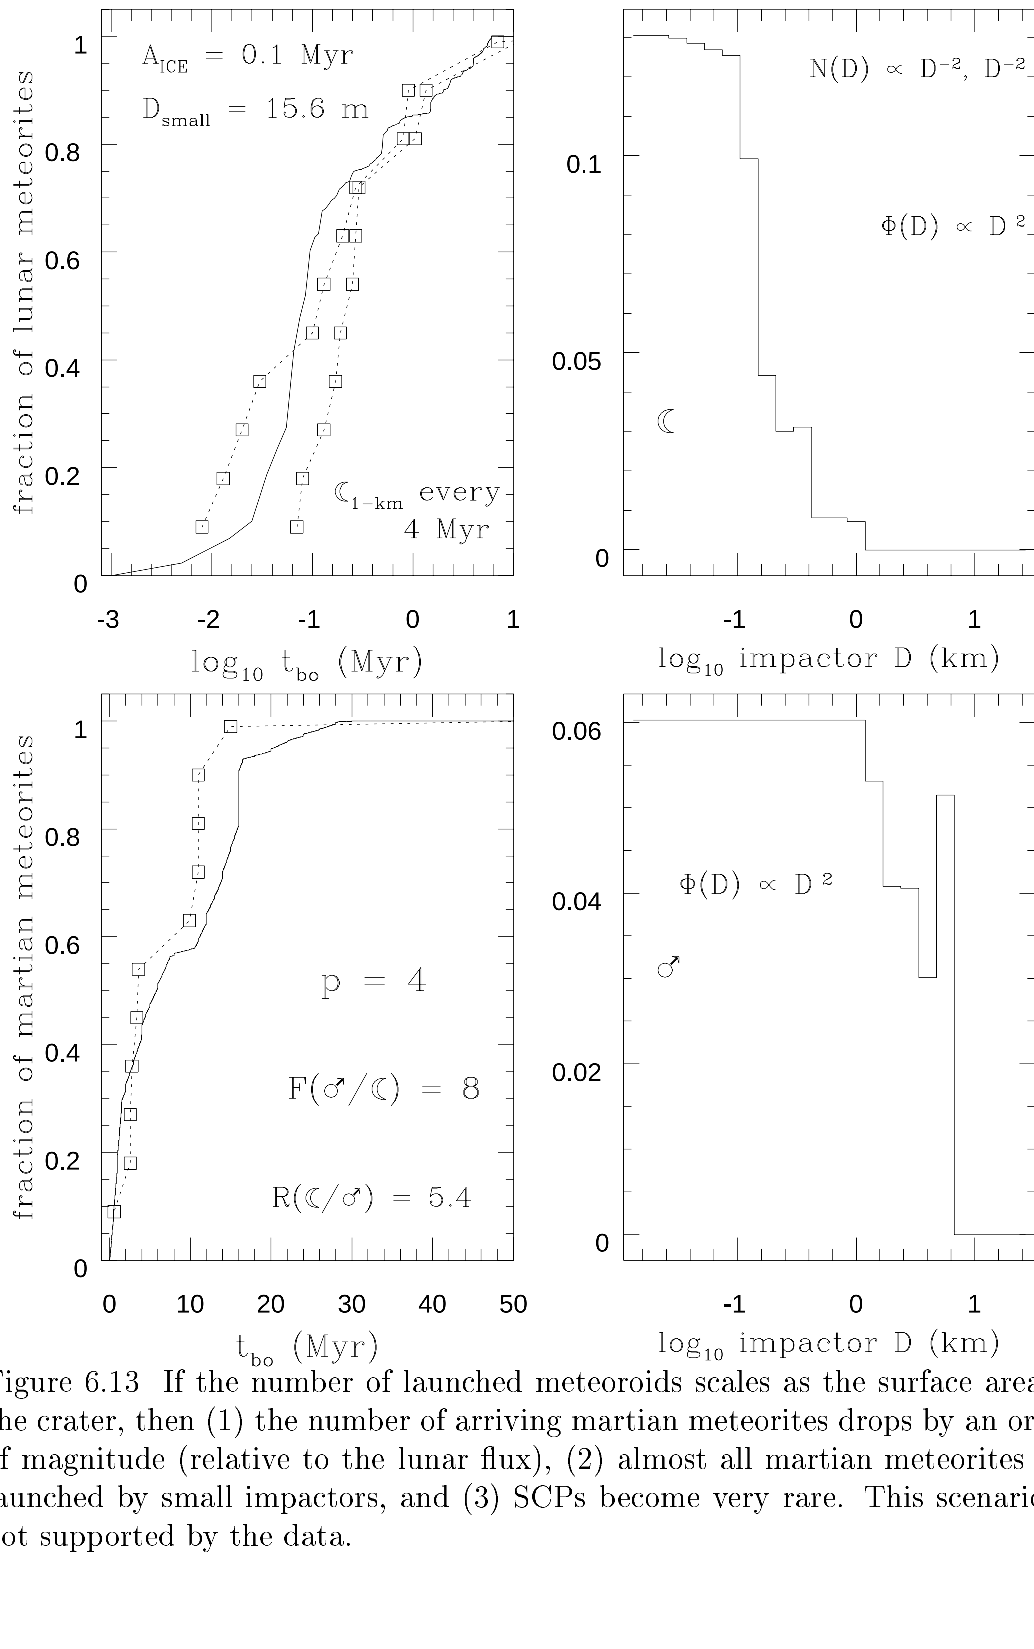

Figure 6.13 If the number of launched meteoroids scales as the surface area of the crater, then (1) the number of arriving martian meteorites drops by an order of magnitude (relative to the lunar flux), (2) almost all martian meteorites are launched by small impactors, and (3) SCPs become very rare. This scenario is not supported by the data.

6.3.3.1 Implications

These results suggest that there must be some process which discriminates against meteoroid production by smaller martian impactors, and yet still allows for large impactors to produce enough of a flux to keep $R \sim 1$. One can imagine several hypotheses to explain this.

(1) Perhaps the martian atmosphere appreciably modifies the size distribution of the impactors or launched meteoroids. It is unlikely that the 6 mbar atmosphere can prevent impactors from reaching the surface; even Earth's more massive atmosphere only stops silicate impactors below $D \sim 100$ m (Chyba 1993). The martian atmosphere will only stop yet smaller impactors, which probably would lack sufficient kinetic energy to liberate meteoroids in any case. On the other hand, the atmosphere may hinder small (up to tens of cm) launched meteoroids from escaping back out to space, since they have much smaller ratios of surface area to mass (the important parameter for drag). Therefore it is conceivable that smaller ejected particles are preferentially prevented from escaping the planet. This would help explain the larger average mass of the martian meteorites.

(2) An entirely different launch process may dominate on Mars. Perhaps the vaporization of subsurface ice could assist meteoroid launch. Any such proposed process must still satisfy all of the constraints (preferences for deep launch, young source terrains, and bigger impactors). The data do not yet seem to demand this extreme hypothesis.

(3) As shown above, if the number of meteoroids scales as the volume of the crater, this naturally resolves the problem. Volume scaling might occur in two ways.

(i) Perhaps the martian surface is not covered with a regolith layer having a sound speed as low as the lunar one, at least not in the young source terrains of most of the martian meteorites. This would result in the martian impactors mining down deeper than their lunar counterparts and launching more ejecta. However, in this case we might then expect many objects from small impactors and thus fewer source-crater pairs.

(ii) Alternately, perhaps the surface layers on the two parent bodies are identical. Then small martian impactors only excavate very shallow surface layers of little volume. But the transfer dynamics show that ancient ($\sim$1-10 Myr ago) large martian impacts are well represented in the impact record. Even *with* a 100 m/s sound speed, a 3-km impactor (not an uncommon occurrence on Mars over 10 Myr) will have a spall zone of depth $z \approx 30$ m (Eq. 6.7). Once down into bedrock, as such an impact would almost surely be, the sound speed would rise dramatically and thus the spall layer may be considerably deeper than 30 m. In this regime the amount of ejecta may become proportional to the volume of the projectile (Melosh 1985, Vickery and Melosh 1987) and almost all the meteorites will come from larger craters (*cf.* launcher spectrum in Fig. 6.12). The scenario has the advantage that it does not matter if the smaller impactors are being screened out or not. The larger masses of the martian meteorites may then be simply due to their being produced in larger impacts; on the other hand, atmospheric effects may still enhance this effect by selecting against smaller meteoroids. The deep provenance of the martian meteorites may be either (1) due to a lack of near-surface rocks suitable for ejection (Warren 1994), or (2) simply because only a very small fraction of the material spalled off in larger impacts is within $\approx$3 m of the surface when depths down to many tens of meters are excavated.

6.3.3.2 The fraction of young terrain

Since 11 of the 12 martian meteorites are from geologically young terrain which occupies a small portion of the surface, this strongly argues that these terrains are favored (perhaps almost exclusively) for launch. Conceivably, small number statistics could be to blame, but if six impact events are in fact represented (Eugster *et al.* 1996), then this hypothesis begins to stretch the bounds of credibility. I have explored including the effect of having limited source regions by simply randomly rejecting 80% of all martian impacts, hypothesizing that only 15% of the martian surface contains sufficiently young terrain (Tremain 1995a), and allowing 5% of the older terrain to be suitable (to allow for the existence of ALH84001). Although clearly ad hoc, this assumption does provide a mechanism for the selection of younger meteorites. It simultaneously increases the frequency of source-crater pairings, since the martian launch process becomes even more dependent on the stochastic nature of the largest impacts, which must now strike inside the 15 Myr 'delivery window', but also into young terrain. Even the baseline case seems to have too little source-crater pairing in the expected spectrum (Fig. 6.11). After reinstating random impactor-timing, allowing only 20% of the martian surface to launch meteorites produces expected spectra such as the example of Fig. 6.14. Here, three large impactors hitting Mars in the last 6 Myr have produced 90% of the martian meteorites; so many meteoroids were in fact produced that just as many martian meteorites were delivered to the ice sheet as lunar ones. Such an outcome is not commonplace however. In most simulations that use the parameters of Fig. 6.14 lunar meteorites are favored by a factor of 3–6; this discrepancy could be easily rectified by minor changes in the relative flux or power law exponent. Of course, it need not be rectified at all; Fig. 6.14 shows that the observed spectra can be reproduced with this set of parameters.

A possibility worthy of mention is that the martian atmosphere is in large part responsible for the selection of young terrains. This could occur because the young volcanic terrains on Mars (specifically in the Tharsis region) are at higher altitudes than the older terrains. If the martian atmosphere does effectively hinder meteoroid launch, as discussed in the previous section, then the higher altitude regions could be favored since considerably less atmospheric mass will

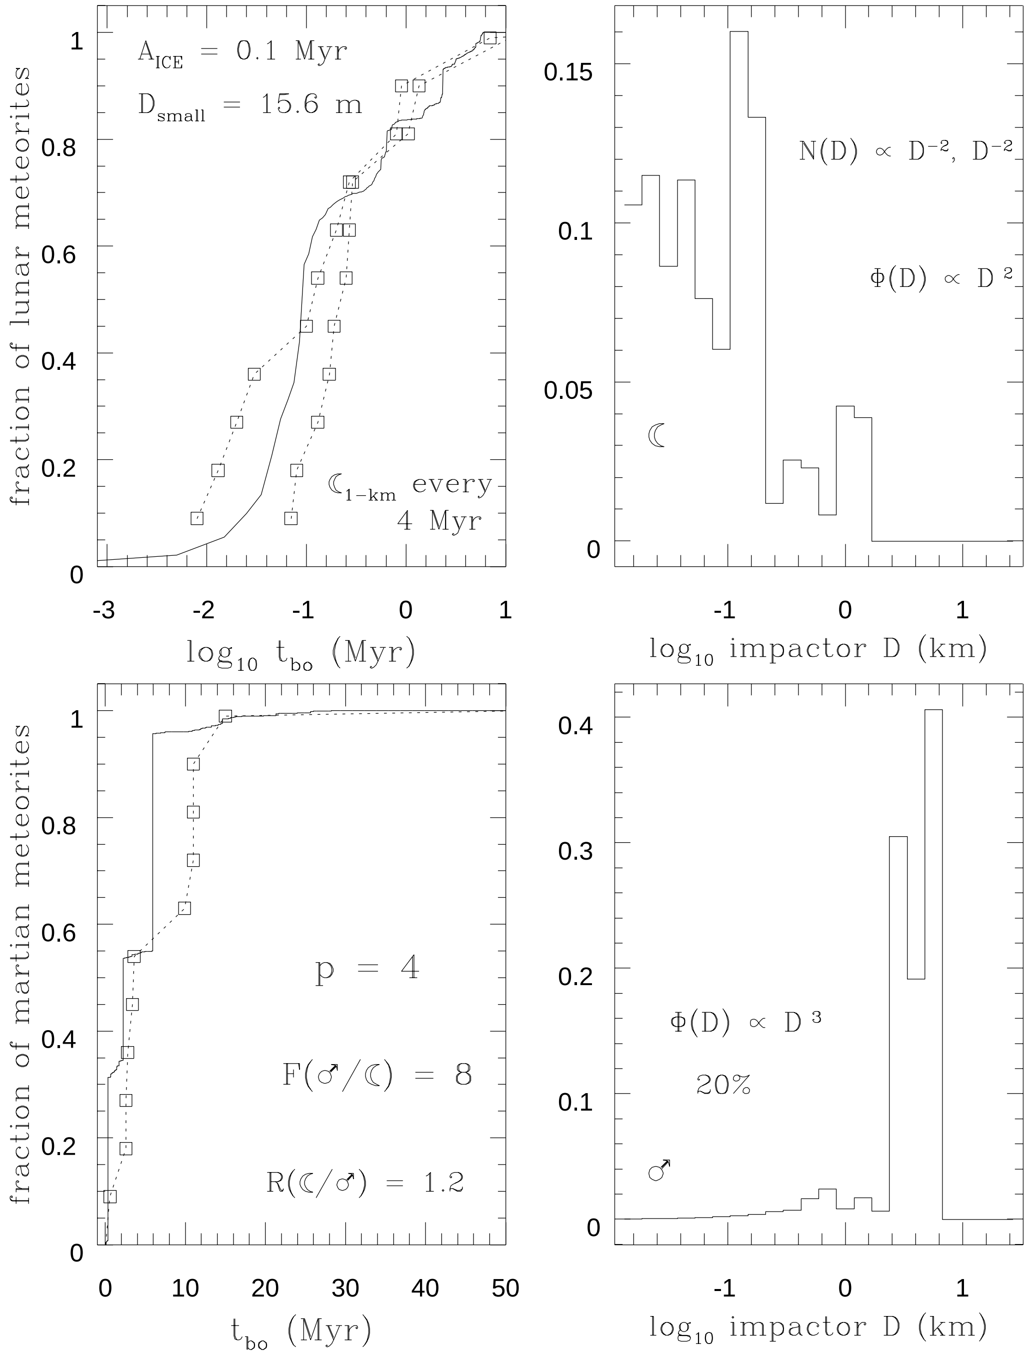

Figure 6.14 Limiting the fraction of Martian terrain results in a greater tendency for source-crater pairing. Although this on average would be expected to decrease the number of launched martian meteoroids, large impacts into young terrain can still produce more than enough meteorites, as in this example where just as many martian as lunar meteorites are present in the ice sheet. The exact form of the martian age spectrum clearly depends on the timing of these impacts.

need to to penetrated by the ejecta. At the 6 mbar level (the zero-altitude reference level on Mars), the amount of atmospheric mass per unit area above the surface is $\sigma = P_\circ/g \simeq 16$ g/cm^2, where $P_\circ$ is the surface pressure and $g =$ 371 cm/s^2 the gravitational acceleration. Meteorite launch would be hindered if the ejected meteoroid needs to traverse roughly its own mass in atmosphere (Chyba 1993). At 0-km altitude, only very large meteoroids ($r \gtrsim 15$ cm, $m \gtrsim$ 30 kg) have ratios of mass-to-surface area larger than necessary. In contrast, at elevations above one martian scale-height (≈ 11 km), 5-cm meteoroids ($m \sim 1$ kg) could escape. The ancient northern highlands have an average elevation of about 4 km (Esposito *et al.* 1992), so only the very highest terrains will be above a scale-height, and thus favored by this atmospheric effect. While this scenario clearly favors launch from the very young volcanic terrains, it further restricts the available portions of the martian surface. In fact, only the flanks of the four largest Tharsis volcanos are above 11 km altitude; Jeffrey E. Moersch (1996, private communication) estimates this represents only $\sim$ 0.3% of the martian surface area.

6.3.3.3 Stochastic variability

The example of Fig. 6.14 illustrates that the stochastic nature of the martian (and sometimes lunar) impacts could be crucial to understanding features in the data. For this reason, to provide the reader with some sense of how much random variation occurs, I show the results of several simulations. Figs. 6.15–6.23, along with Fig. 6.14, present the outcomes of ten sequential simulations, all with the same parameters. The differences in these plots are entirely due to the random-impact timing and the limited viable martian launch area. These figures demonstrate that: (1) the exact nature of the martian age spectrum in regards to source-crater pairings depends heavily on the timing of the largest impacts, and (2) a lunar-to-martian meteorite number ratio $R \sim 1$ is not in fact uncommon; the ratio is usually within a factor of 6 of this.

Fig. 6.15 is an especially interesting case, because it illustrates the consequences of a large, recent lunar impact. 100 Kyr after the event, the ice sheet is still so dominated by this 2-km impactor's ejecta that almost all lunar meteorites come from the impact and these meteorites outnumber the martin ones by a factor of 100. Only after the hail of meteorites from this impact ceases (after a few hundred kyr at most), and are destroyed in the ice sheet, does R return to order 1–6, as it is most of the time.

6.3.4 Musings on martian source-crater pairings

An unresolved issue in the previous discussion is the extent of source-crater pairing (SCP) in the martian meteorites. As discussed in Sec. 4.4, it is almost inconceivable in light of the constraints that all of the martian meteorites were

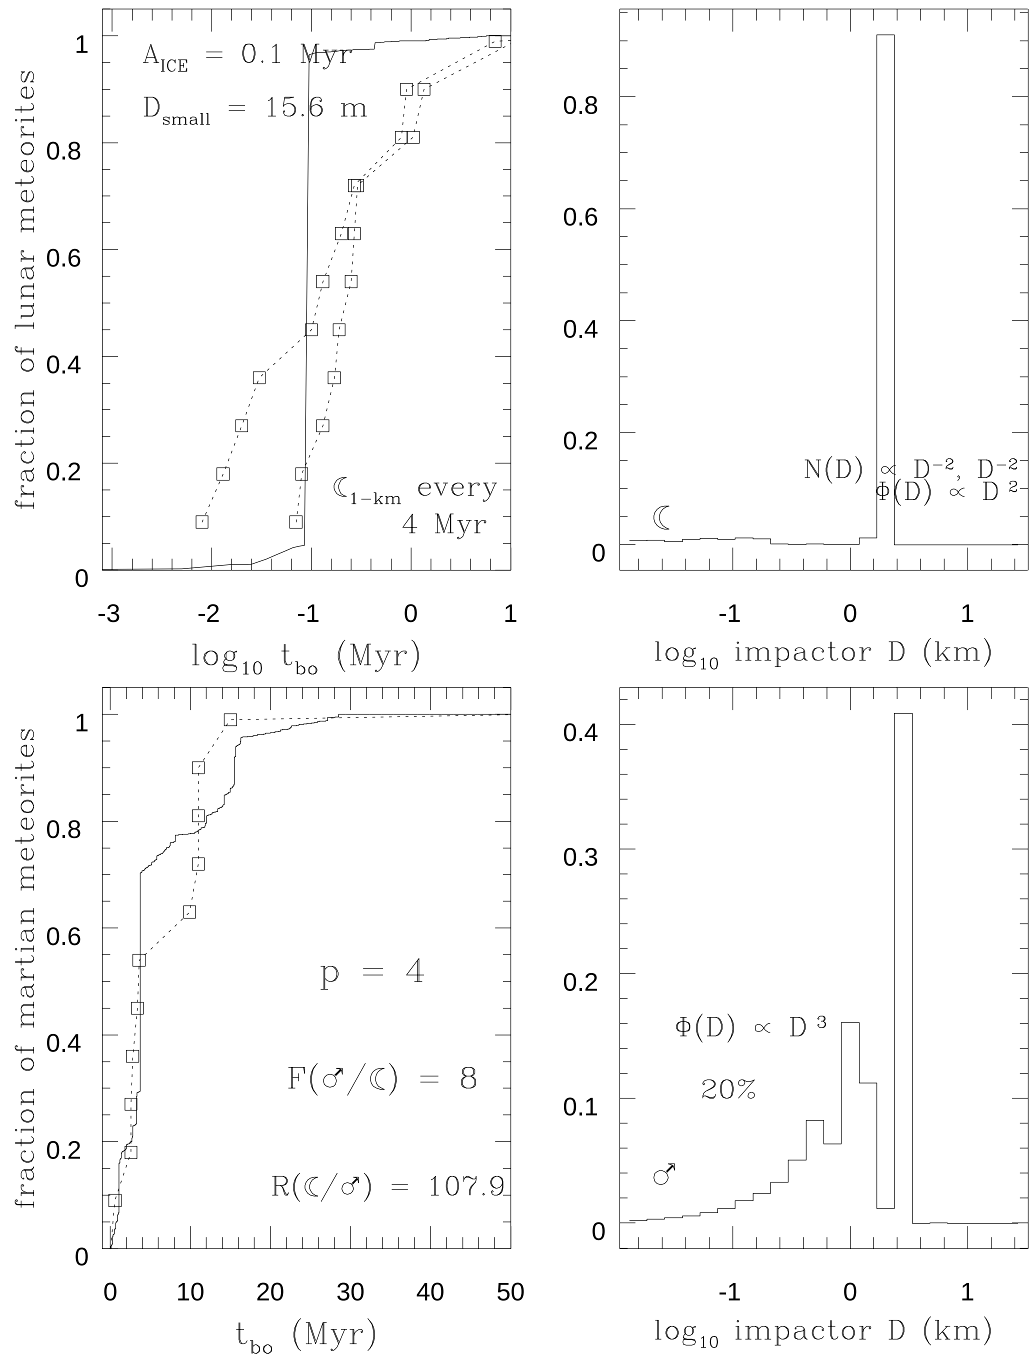

Figure 6.15 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

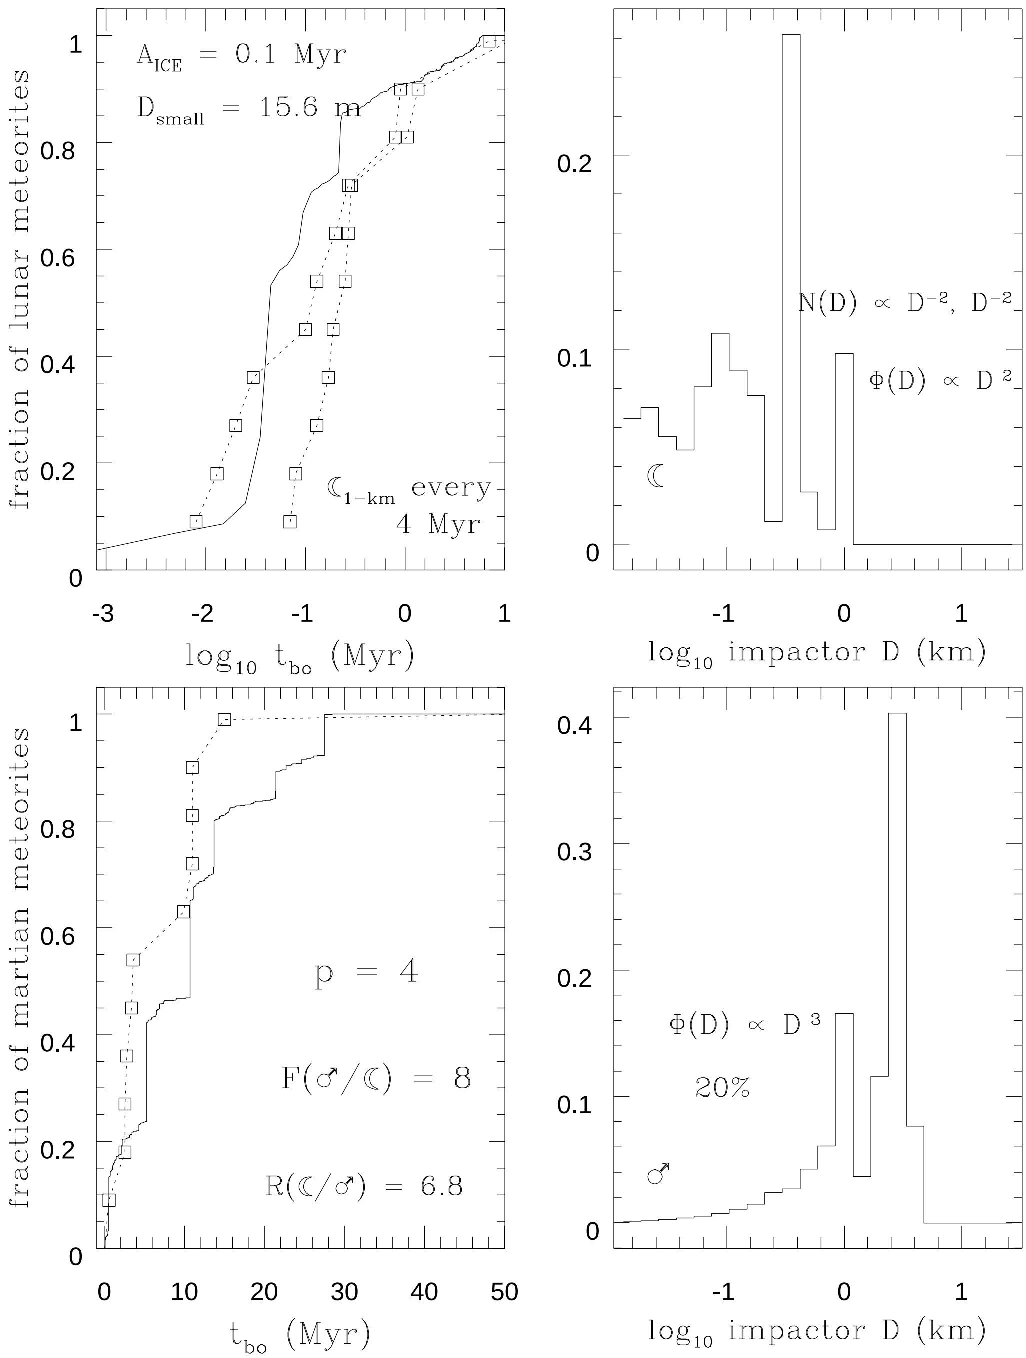

Figure 6.16 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

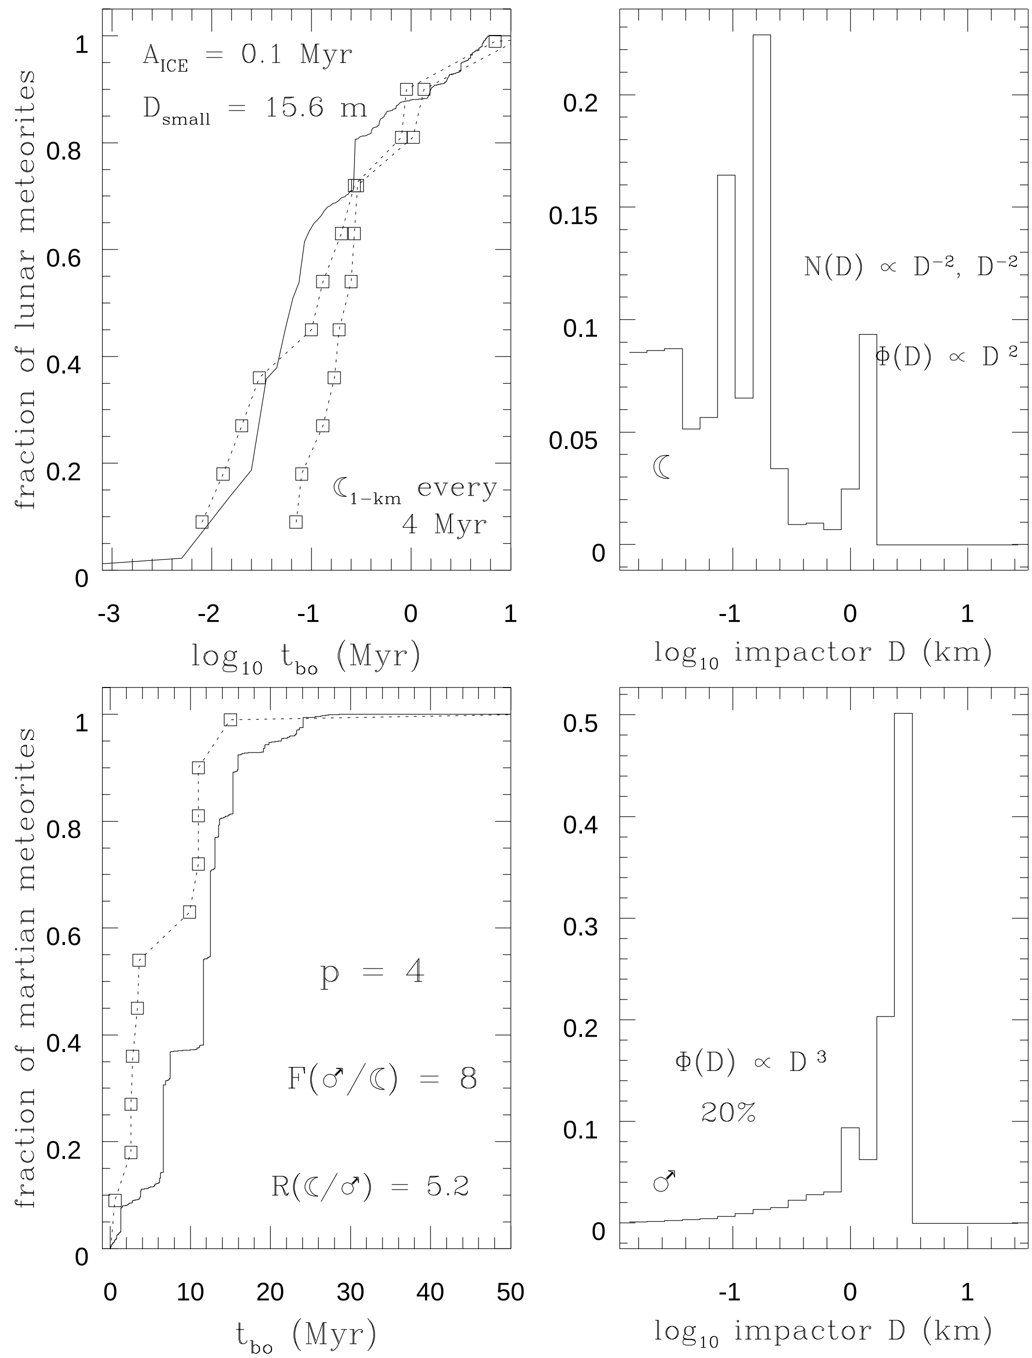

Figure 6.17 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

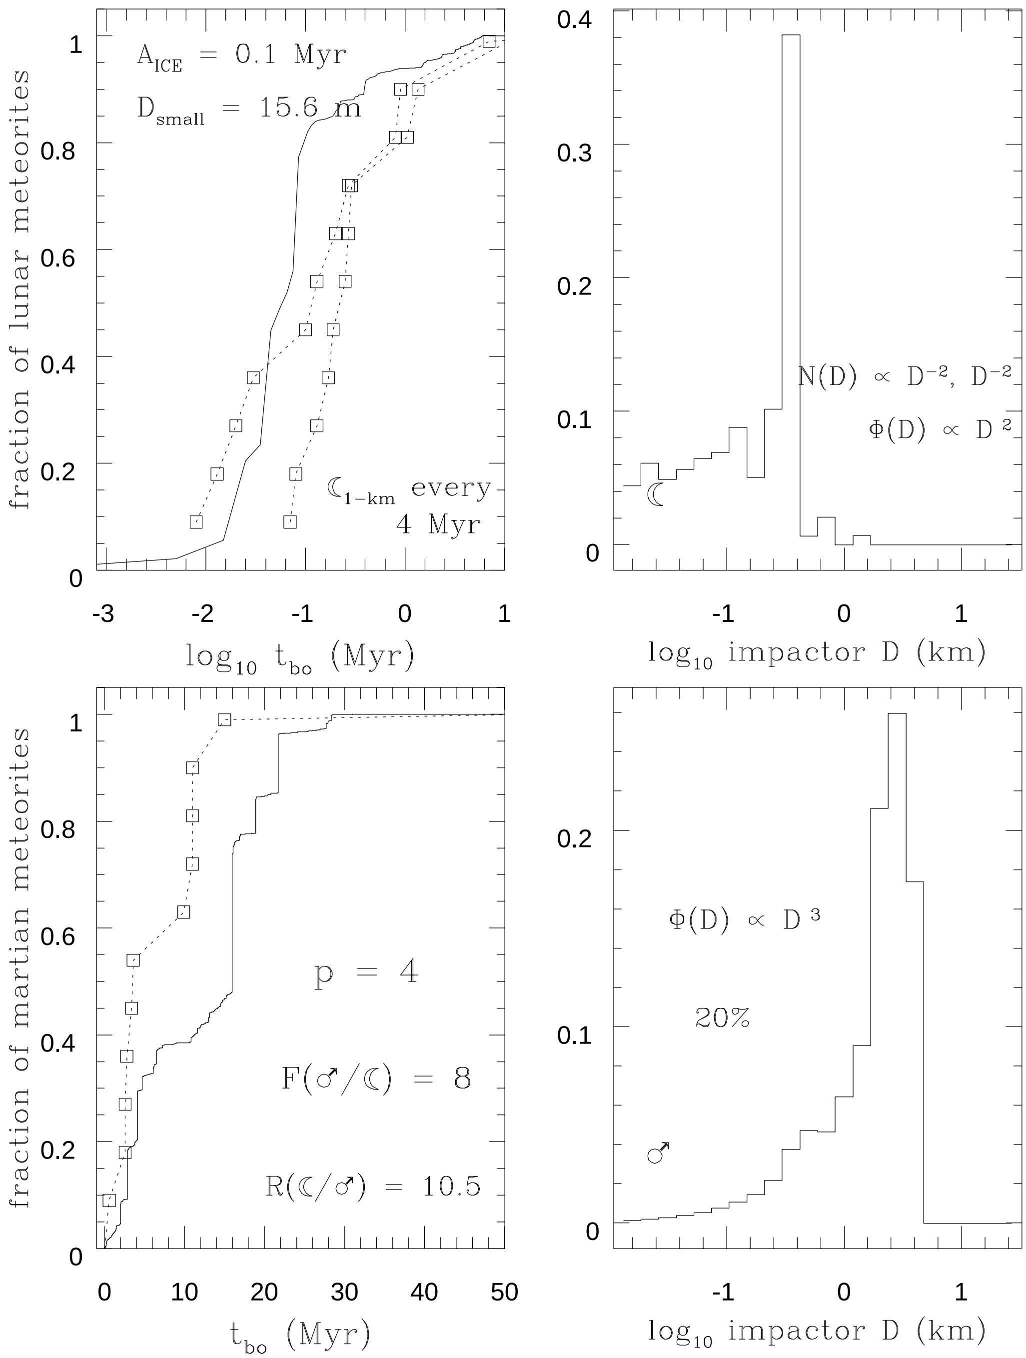

Figure 6.18 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

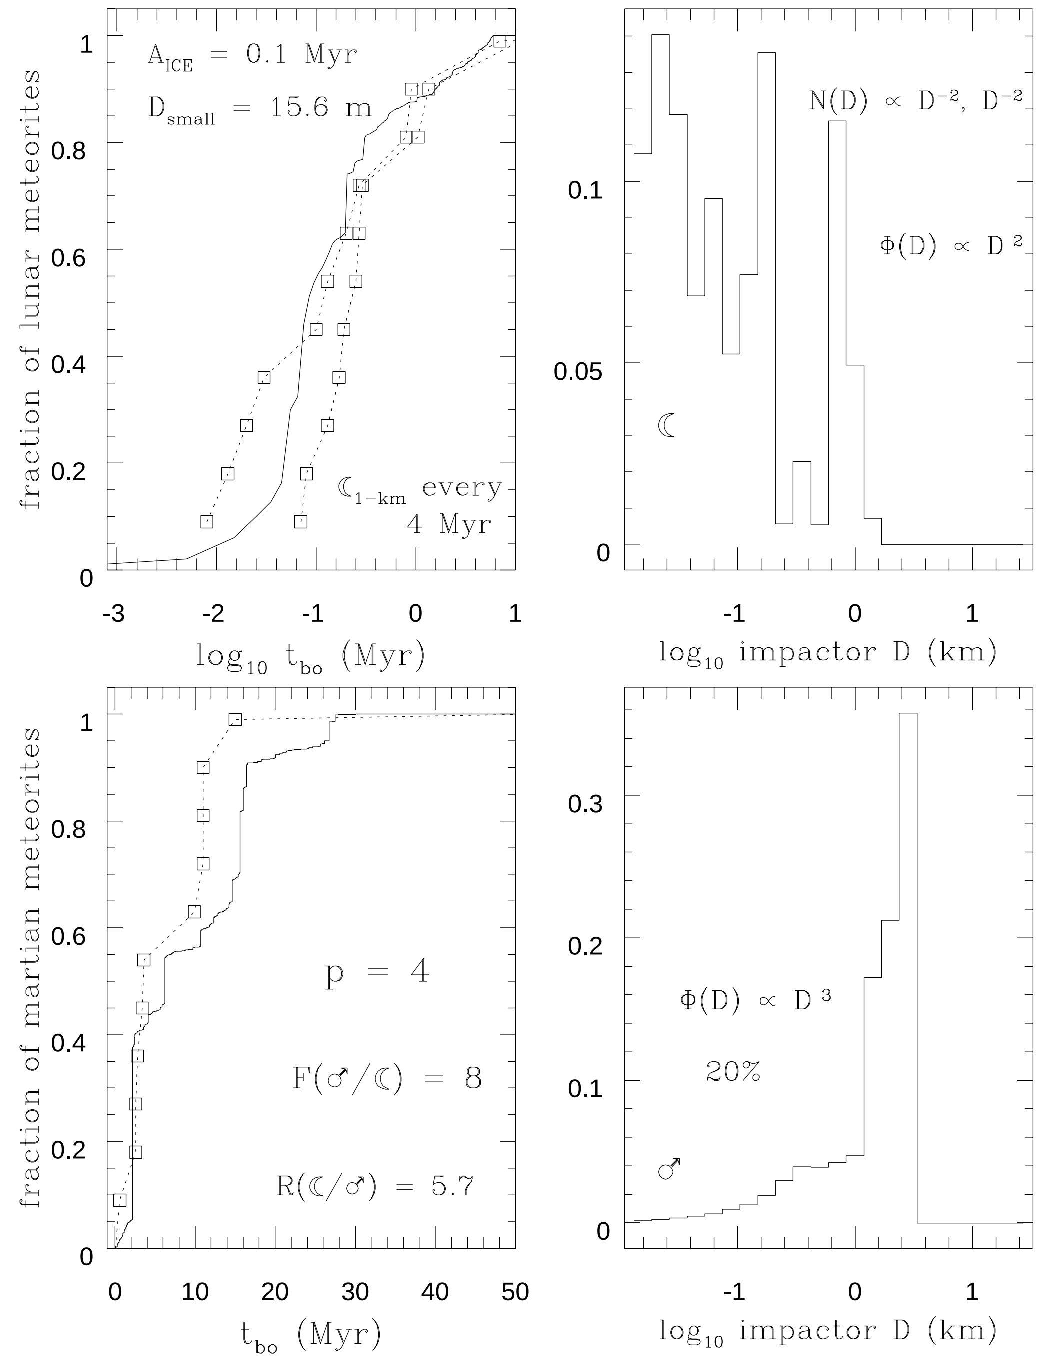

Figure 6.19 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

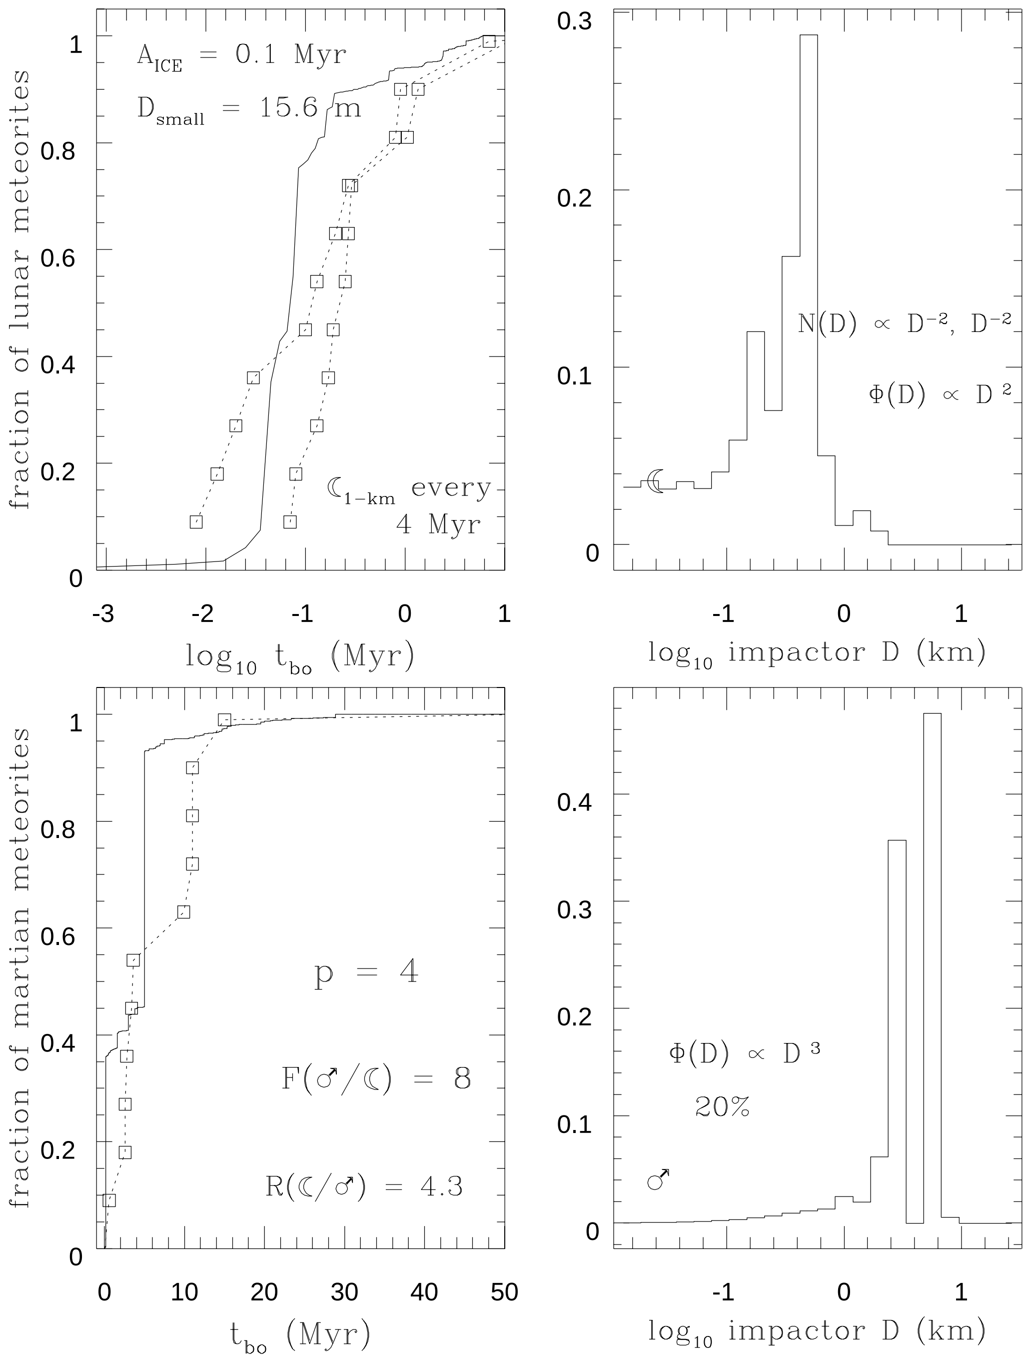

Figure 6.20 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

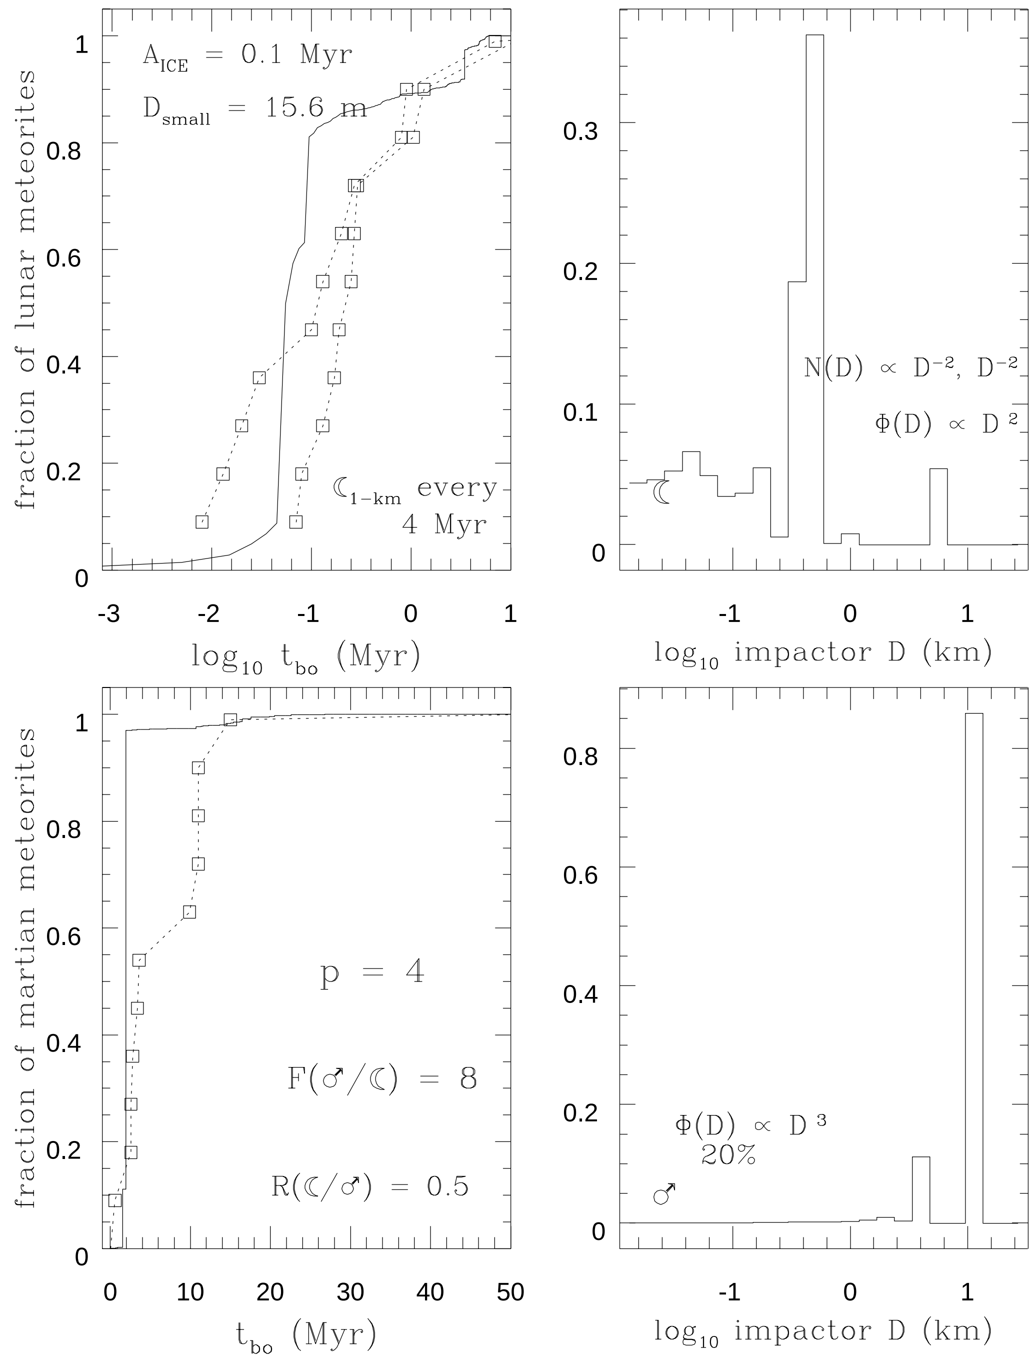

Figure 6.21 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

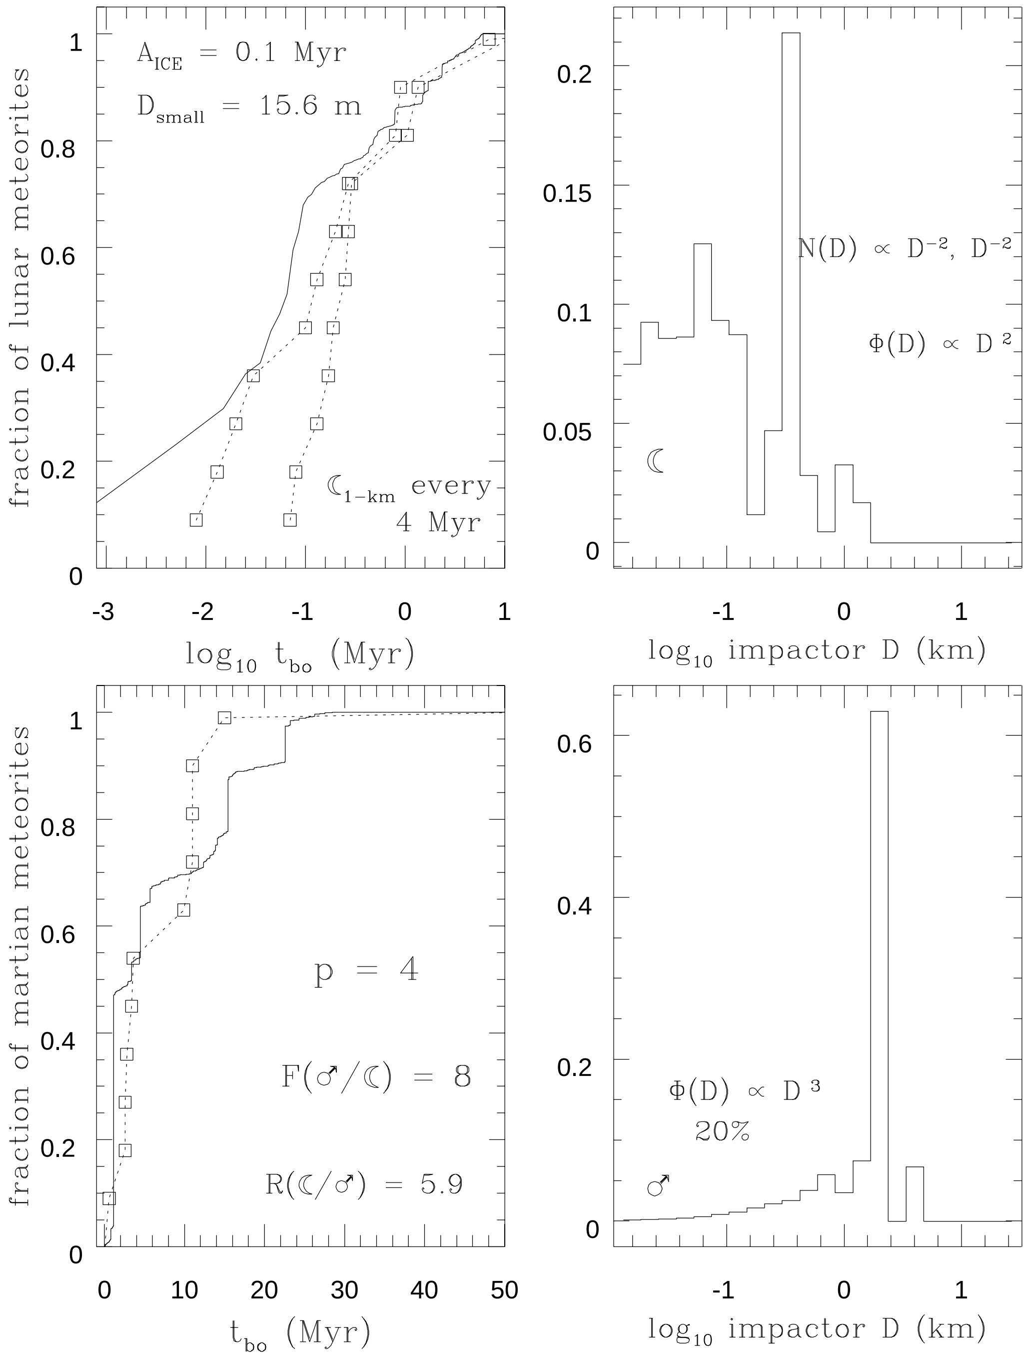

Figure 6.22 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

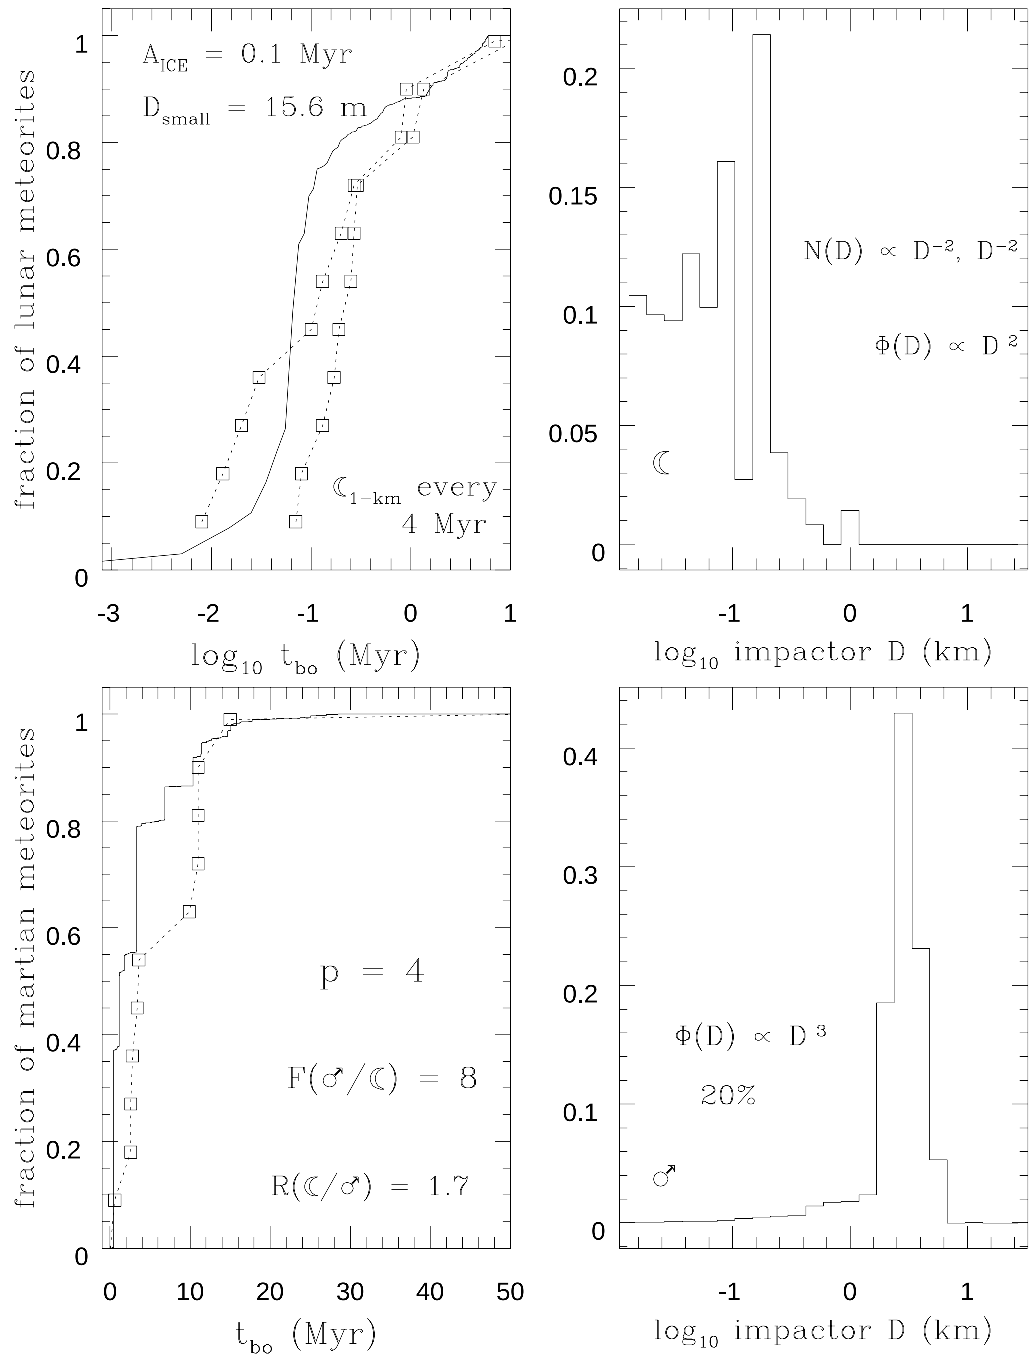

Figure 6.23 A sample delivery history for lunar and martian meteorites. This is part of a series of ten simulations (Figs. 6.14–6.23), all of which use the same parameters, intended to illustrate the importance of large stochastic impacts on the Moon and Mars. These simulations use randomly timed impacts and limited martian source area.

launched in a single impact. However, one could attempt to make a case for only three impacts being represented: ALH84001 15 Myr ago, the NCs about 10 Myr ago, and all the shergottites about 3 Myr ago. This scenario would require the 10-Myr impact to overlap both N and C source terrains. It would also require the 3-Myr impact to launch both lherzolitic and basaltic lithologies *and* to produce very large meteoroids in order to shield at least EET79001 for 2 Myr (and perhaps the basaltic shergottites for ≈1 Myr). Such a hypothesis eases the problem involving the young source terrains, for then only 2 of 3 impacts are into the young regions instead of needing 5 of 6 to be so. This hypothesis thus demands a very large 3-Myr impact, which is plausible as indicated by the variety of possible histories seen in Figs. 6.14–6.23.

The model presented in this chapter is unable, given the amount of data, to distinguish between the likelihood of having 3 versus 6 martian impacts represented, although I prefer the latter option. It is unlikely that EET79001 is a shielded collisional fragment from a 2.5-Myr event *and* that the lherzolitic shergottites are from a separate impact event; I believe it to be all or nothing (*i.e.,* all the shergottites were launched in a single event with partial shielding, or three events occurred).

These questions are likely unanswerable at present. In fact, it may be fair to say that they may never be answered.

6.4 The big picture

I now summarize the general theory, and put forward a hypothesis that resolves the puzzles presented in the introduction regarding the lunar and martian meteorites. The data are not yet plentiful enough to constrain these conjectures with certainty; the following general picture does explain the available facts, but perhaps not uniquely.

The lunar and martian surfaces are bombarded by a set of impactors whose diameters fall off as D^{-2}. These impactors launch meteoroids by stress-wave interference, from a spall zone of depth $z \sim Dc_L/U$. In the Moon's surface layers the sound speed $c_L \simeq$100 m/s so that impactors of tens to hundreds of meters in diameter (striking with $U \sim$10 km/s), can only loft into space a thin spall layer of a few meters depth. Impactors large enough to tunnel below the lunar regolith do not tend to strike often enough be represented in the present-day meteorite sample. In contrast, such large impactors ($D \gtrsim 1$ km) do hit Mars on time scales less than 10 Myr, where these impactors launch ejecta in greater quantity and from greater depth.

Once ejected into space, 40±15% of lunar meteoroids find their way to Earth, whereas 4±2% of martian meteoroids do. The lunar meteorites mostly fall in a fairly brief burst immediately after launch; the majority arrive in 100,000 years, with only a relatively slow trickle continuing after a million years. In contrast, the martian meteorites are delivered in an almost steady stream over about 10–20 Myr, where the upper limit is caused by a combination of collisional destruction and orbital loss. On this basis alone we expect the large impacts either represented in the ice sheet or recovered as falls to be solely martian. A slightly oversimplified statement would be that the lunar meteorites we have are launched by 10–300 m impactors that strike the Moon every tens of thousands to few hundred thousand years, while the martian meteorites we have were launched by 1–10 km impactors that are striking Mars every 1–10 Myr.

I propose the following resolution of the 5 puzzles listed on p. 7, which I present in the reverse of the order given there.

(5) The martian age spectrum terminates at ≈15 Myr because Sun-grazing and collisional destruction eliminate the meteoroid population in space on this time scale.

(4) The prevalence of geologically young rocks amongst the martian meteorites is either a statistical fluke, due to pervasive weathering which makes the older regions of Mars unsuitable for meteorite launch (as suggested by Warren 1994), or caused by martian atmospheric effects. This is probably the most important unresolved puzzle.

(3) The lack of shallow-provenance martian meteorites comes about simply because the upper 3 m of surface material comprises a tiny fraction of the very thick spall layers thrown off in the large impacts that are responsible for launching martian meteoroids.

(2) The martian meteorites have larger mean masses because they are produced by larger impactors. The martian atmosphere may play some role in preventing smaller meteoroids from escaping.

(1) The equal numbers of lunar and martian meteorites is a consequence of the finite age of the ice sheet combined with transfer dynamics that today deliver few meteorites from ancient lunar impacts, but a reasonable flux from larger ancient martian impacts.