

On the red supergiant problem

C. S. Kochanek^{1,2}★

¹Department of Astronomy, The Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA

²Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 W. Woodruff Avenue, Columbus, OH 43210, USA

Accepted 2020 February 22. Received 2020 February 20; in original form 2020 January 20

ABSTRACT

We examine the problem of estimating the mass range corresponding to the observed red supergiant (RSG) progenitors of Type IIP supernovae. Using Monte Carlo simulations designed to reproduce the properties of the observations, we find that the approach of Davies & Beasor significantly overestimates the maximum mass, yielding an upper limit of $M_h/M_\odot = 20.5 \pm 2.6$ for an input population with $M_h/M_\odot = 18$. Our preferred Bayesian approach does better, with $M_h/M_\odot = 18.6 \pm 2.1$ for the same input populations, but also tends to overestimate M_h . For the actual progenitor sample and a Salpeter initial mass function, we find $M_h/M_\odot = 19.01_{-2.04}^{+4.04}$ for the Eldridge & Tout mass–luminosity relation used by Smartt and Davies & Beasor, and $M_h/M_\odot = 21.28_{-2.28}^{+4.52}$ for the Sukhbold, Woosley & Heger mass–luminosity relation. Based on the Monte Carlo simulations, we estimate that these are overestimated by $(3.3 \pm 0.8) M_\odot$. The red supergiant problem remains.

Key words: stars: massive – supernovae: general.

1 INTRODUCTION

Particularly as the archive of *Hubble Space Telescope* images of nearby galaxies has grown, there has been steady progress in identifying the progenitors of core-collapse supernovae (ccSNe, see the reviews by Smartt 2015; Smartt 2009). In Kochanek et al. (2008), we pointed out that there appeared to be a deficit of higher mass progenitor stars. This point was made more clear with better statistics for Type IIP ccSNe by Smartt et al. (2009). The progenitors of Type IIP ccSNe are red supergiants (RSGs), and Smartt et al. (2009) found progenitors with masses between $8.5_{-1.5}^{+1.0}$ and $(16.5 \pm 1.5) M_\odot$, while RSGs in the Local Group are found with masses of up to $25 M_\odot$. Smartt et al. (2009) termed this the ‘red supergiant problem’. There is also evidence to support the existence of a problem from studies of stellar populations near SNe or SN remnants (see, e.g. Díaz-Rodríguez et al. 2018; Auchettl et al. 2019) and the lack of the X-ray bright SNe that would be expected given the higher mass-loss rates of the more massive progenitors (Dwarkadas 2014).

The red supergiant problem could be solved by eliminating the gap between the highest observed progenitor masses and the predicted maximum masses at which stars explode as RSGs. One possibility is to modify stellar evolution and mass-loss to reduce the maximum mass of stars exploding as RSGs and have them instead explode as Type Ib or Type Ic ccSNe (Groh et al. 2013). To be a solution, they need to explode as stripped stars, not blue supergiants. Almost no SN light curves are consistent with the explosions of BSGs (1 in 92 in Smartt 2009) and no BSG progenitors (other than

SN 1987A) have been identified even though they should be no more difficult to be identified than RSG progenitors.

A second possibility is to posit that the luminosities, and hence the masses, of the observed progenitors have been underestimated due to unrecognized foreground or circumstellar extinction (Walmswell & Eldridge 2012; Beasor & Davies 2016; Davies & Beasor 2018). Note, however, that it is easy to overestimate the effects of dust (Kochanek, Khan & Dai 2012). A third possibility is to argue that the difference between the maximum masses of progenitors and the expected maximum masses of RSGs is statistically insignificant (Davies & Beasor 2018). Zapartas et al. (2020) find that the effects of binary evolution have essentially no consequences for the red supergiant problem.

The alternative physical explanation is that the missing progenitors are not being found because the more massive RSGs are not exploding as SNe and instead become black holes (Kochanek et al. 2008). Stars in the mass range of the missing RSG progenitors have internal structures that are particularly difficult to explode (O’Connor & Ott 2011; Ugliano et al. 2012; Pejcha & Thompson 2015; Sukhbold et al. 2016), and failed explosions of these RSGs provide the first natural explanation for the observed masses of Galactic black holes Kochanek (2014, 2015). Furthermore, our search for failed SNe with the Large Binocular Telescope (Gerke, Kochanek & Stanek 2015; Adams et al. 2017a,b) has identified one excellent candidate for a failed ccSN whose estimated progenitor mass is exactly in the range needed to explain the red supergiant problem and the masses of the Galactic black holes. The failed ccSN rate implied by the discovery of one candidate is also consistent with theoretical expectations.

In this paper, we reconsider the problem of estimating the mass range of RSG progenitors. We assume that stars explode in a mass range from M_l to M_h with a Salpeter (1955) power-law initial

* E-mail: ckochanek@astronomy.ohio-state.edu

mass function (IMF), $dn/dM \propto M^{(-1-x)}$ with $x = 1.35$. In modern examinations of the explodability of stars (e.g. O'Connor & Ott 2011; Ugliano et al. 2012; Pejcha & Thompson 2015; Sukhbold et al. 2016), the relationship between mass and outcome is more complex, with explosions and failures interspersed in mass, but there is still effectively a maximum mass. The objective is to estimate the two mass limits M_l and M_h . In particular, Davies & Beasor (2018) carries out an analysis to find an upper limit of $M_h/M_\odot = 19.0^{+2.5}_{-1.3}$, which is significantly above the estimate of $M_h/M_\odot = 16.5 \pm 1.5$ by Smartt (2009). Davies & Beasor (2018) further argue that this should be corrected to $M_h/M_\odot = 25$ because the highest mass observed progenitor must lie below M_h , leading to an underestimate of the limit that requires an upward correction.

A simple way to examine this question is to use Monte Carlo simulations designed to closely mimic the properties of the observations and then analyse them to see how well the input mass limits are recovered. We will consider both the Davies & Beasor (2018) analysis method and a Bayesian approach that is similar in spirit to the original Smartt (2009) analysis. In Section 2, we describe the calculations, and in Section 3, we discuss the results.

2 METHODS

For this paper, we simply adopt the tabulation of the properties of 24 Type II progenitors from Davies & Beasor (2018). The progenitors are characterized by a distance modulus, μ , a broad-band filter magnitude or magnitude limit, m_λ , an estimated extinction for that wavelength, A_λ , and a bolometric correction, BC_λ . Davies & Beasor (2018) treat SN 2009md slightly differently, but we filled in the missing values in their Table 4 so as to reproduce their estimates of the progenitor luminosity and its uncertainties. Associated with each quantity is an uncertainty: σ_μ , σ_m , σ_A , and σ_{BC} . We also require statistical distributions for these quantities. Davies & Beasor (2018) treat the distributions as Gaussians except for the bolometric correction, which is viewed as uniformly distributed between $BC_\lambda - \sigma_{BC}$ and $BC_\lambda + \sigma_{BC}$. Davies & Beasor (2018) round negative extinctions in the tails of the Gaussian extinction distribution upwards to zero. There are 14 flux measurements and 10 upper limits. Where there are flux limits, they are all 3σ limits, with the exception of a 5σ limit for SN 2002hh. As a slight simplification, we convert this into a 3σ limit so that all the limits can be treated uniformly.

Given these quantities, the progenitor luminosity L is

$$2.5 \log \left(\frac{L}{78.6 L_\odot} \right) = -m_\lambda + \mu + A_\lambda - BC_\lambda. \quad (1)$$

The data really determines only a minimum and maximum progenitor luminosity, but this can be converted into a mass range given a mass–luminosity relation. Fig. 1 shows the end-of-life mass–luminosity relations from Schaller et al. (1992), Eldridge & Tout (2004), Groh et al. (2013), and Sukhbold et al. (2018). Smartt (2009) and Davies & Beasor (2018) primarily used the Eldridge & Tout (2004) models after eliminating the luminous asymptotic giant branch (AGB) phase for lower mass stars. Each model has some mass above which the models cease to be RSGs at death. Sukhbold et al. (2018) includes models with their standard mass-loss rate, half that rate and one-tenth of that rate, with the stars remaining as RSGs up to 26, 39, and $60 M_\odot$. In Fig. 1, these three mass-loss sequences are virtually indistinguishable, essentially because the mass of the envelope has no effect on the luminosity of the helium core.

For our calculations, we need a mass–luminosity relation that extends beyond the mass range assumed to explode as an RSG, so we use the low-mass-loss models from Sukhbold et al. (2018)

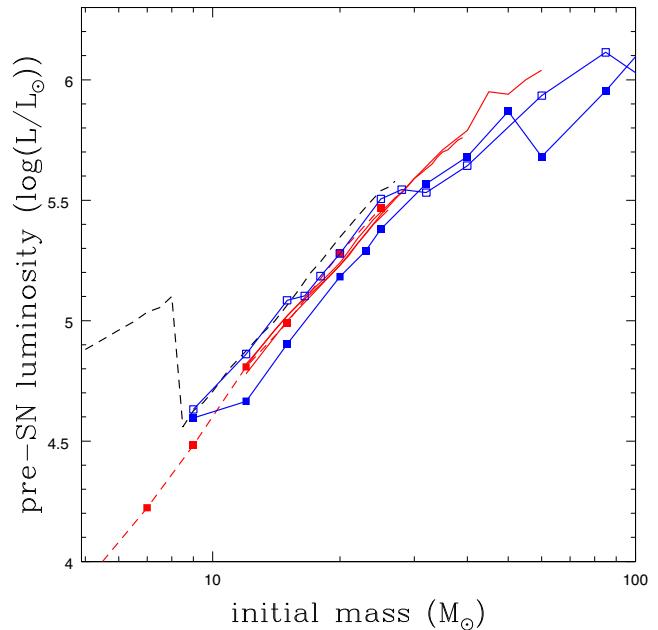


Figure 1. End-of-life mass–luminosity relations from Schaller et al. (1992) (red dotted, filled squares), Eldridge & Tout (2004) (black dashed), non-rotating Groh et al. (2013) (blue solid, filled squares), rotating Groh et al. (2013) (blue solid, open squares), and Sukhbold et al. (2018) (red solid). The three mass-loss models from Sukhbold et al. (2018) lie almost on top of one another. Only the Eldridge & Tout (2004) models include high-luminosity AGB phase at lower masses.

extended to lower mass ($< 12 M_\odot$) using the models from Schaller et al. (1992) since the two sets of models overlap. For ease of calculation,

$$\log \frac{L}{L_\odot} = 4.610 + 2.267 \log \left(\frac{M}{10 M_\odot} \right) - 0.494 \log^2 \left(\frac{M}{10 M_\odot} \right) \quad (2)$$

and

$$\log \frac{M}{M_\odot} = 1.180 + 0.489 \log \left(\frac{L}{10^5 L_\odot} \right) + 0.056 \log^2 \left(\frac{L}{10^5 L_\odot} \right) \quad (3)$$

provide very good polynomial fits to the resulting mass–luminosity relation for $5 < M < 60 M_\odot$. The shape of the polynomials also fits the Eldridge & Tout (2004) models well, but the leading constants become 4.703 and 1.131 for the luminosity and mass, respectively. The offsets mean that the Sukhbold et al. (2018) models are 24 per cent less luminous at fixed mass and 12 per cent more massive at fixed luminosity than the Eldridge & Tout (2004) models.

Like Davies & Beasor (2018), we simply assume a Salpeter (1955) IMF, $dN/dM \propto M^{(-x-1)}$ with $x = 1.35$ leading to an integral distribution of progenitor masses of

$$P_{\text{SN}}(< M) = \frac{M_1^{-x} - M^{-x}}{M_1^{-x} - M_h^{-x}} \quad (4)$$

over the mass range $M_l \leq M \leq M_h$. This can then be inverted to get the mass

$$M_{\text{SN}}(P) = M_l \left[1 - P \left(1 - \left(\frac{M_h}{M_l} \right)^{-x} \right) \right]^{-1/x}, \quad (5)$$

corresponding to a fraction P of the progenitor distribution. The goal is to estimate the two mass limits M_l and M_h given the properties of the progenitors.

We build Monte Carlo test samples similar to the data as follows. First, we estimate the 1σ noise level for each of the 24 SNe. For the flux limits, this simply means dividing the stated flux limit by the stated statistical significance. For those with measurements, we assume, as is almost certainly the case, that the noise is background dominated. We then convert the progenitor magnitude and its error into a flux and its error, and the flux error should correspond to the 1σ noise level of the data. Next, we assume a minimum and a maximum mass, where we used $M_l = 8 M_\odot$ and $M_h = 18 M_\odot$ or $21 M_\odot$, and then randomly draw a mass using equation (5) for each SN. This provides a luminosity through equation (2), which we convert into an apparent magnitude by randomly drawing a distance modulus (Gaussian), extinction (Gaussian rounded up to zero), and a bolometric correction (uniform) using equation (1). If the resulting magnitude is above 3σ , it is treated as a measurement and if it is below, we use the flux limit instead. This produces a random sample of progenitors and limits with the statistical properties of the data.

Davies & Beasor (2018) make 10^5 Monte Carlo trials to estimate the mass limits. For each trial, they randomly draw distances (Gaussian), magnitudes (Gaussian), extinctions (Gaussian rounded up to zero), and bolometric corrections (uniform) for each SN i to derive a luminosity L_i , which is then converted into a mass M_i . For the progenitors with only upper flux limits, the magnitude is taken to be the stated limit, leading to an upper limit on the luminosity and mass for the progenitor in the trial. They then sort the masses and mass limits, discarding any upper mass limits above the highest mass measurement, to leave N' objects. If we index these objects as $j = 0$ to $N' - 1$ and define $u_i = 1$ for detections and $u_i = 0$ for non-detections, they minimize the statistic

$$\chi^2 = \sum_{j=0}^{N'} u_j \left[M_j - M_{\text{SN}} \left(\frac{j}{N' - 1} \right) \right]^2 \quad (6)$$

to estimate M_l and M_h . Note that only the detections ($u_j \equiv 1$) contribute to the statistic, with the highest mass detection having $M_{\text{SN}} = M_h$. The lowest mass detection or upper limit has $M_{\text{SN}} = M_l$. The distribution of the resulting 10^5 values of the M_l and M_h that minimize this fit statistic for each Monte Carlo trial provides their estimate of the allowed minimum and maximum progenitor masses. Since the masses are just weighted uniformly in the χ^2 , and only the maximum likelihood estimates of M_l and M_h are used from each trial, there is no need to define the usual error term in the denominator of the χ^2 .

We prefer a more Bayesian approach that is similar to the original procedures of Smartt et al. (2009), although we will keep the same input data and the relations between fluxes and luminosities as used by Davies & Beasor (2018). We first construct the relative probability distribution that progenitor i has mass M given the data d . For a source with a detection, we compute

$$P_i(M_i|d) \propto \int d\mu dA_\lambda dBC_\lambda P(\mu)P(A_\lambda)P(BC_\lambda) \times \exp \left(-\frac{(m_{\text{mod}} - m_\lambda)^2}{2\sigma_m^2} \right), \quad (7)$$

where the model magnitude m_{mod} comes from rearranging equation (1). For the upper limits, we compute the probability given the

mass that the flux would not exceed the 3σ flux limit,

$$P_i(M_i|d) \propto \int d\mu dA_\lambda dBC_\lambda P(\mu)P(A_\lambda)P(BC_\lambda) \times \text{Erfc} \left(\frac{F_{\text{mod}} - 3\sigma_F}{\sqrt{2}\sigma_F} \right), \quad (8)$$

where the various magnitudes must be converted into fluxes and $\text{Erfc}(x)$ is the complementary error function. $P(\mu)$, $P(A_\lambda)$ and $P(BC_\lambda)$ are the same probability distributions as were used above. We do not need the normalizations of these probability distributions.

Next we must compute the probability of these mass probability distributions given the progenitor mass function. We again use the same fixed relationships between mass and luminosity. For each progenitor, we need to marginalize over the luminosity to get

$$P_i(M_l, M_h|d) \propto \int_{M_l}^{M_h} dM_i P_i(M_i|d)P(M_i|M_l, M_h), \quad (9)$$

where $P(M_i|M_l, M_h) = dP(< M)/dM$ is the probability of having mass M_i given M_l and M_h derived from equation (4). Note that we are maximizing the probability of the detections having their observed fluxes, and the probability that the non-detections are not detected. The final probability distribution for the parameters of the mass function is then

$$P(M_l, M_h|d) \propto P(M_l)P(M_h)\Pi_i P_i(M_i, M_h|d), \quad (10)$$

although in practice we compute $\log P(M_l, M_h|d)$ to avoid floating point underflow problems. We use standard logarithmic priors for the mass limits, with $P(M_l) \propto M_l^{-1}$ and $P(M_h) \propto M_h^{-1}$. The final distribution is normalized so that $\int P(M_l, M_h|d)dM_l dM_h \equiv 1$, and the distribution for one mass limit is found by projecting out the other (i.e. $P(M_l|d) = \int dM_h P(M_l, M_h|d)$).

3 RESULTS AND DISCUSSION

To compare these two approaches, we generated 500 simulated progenitor data sets with minimum and maximum masses of $M_l = 8 M_\odot$ and $M_h = 18 M_\odot$, and then analysed them using either the approach of Davies & Beasor (2018) or the Bayesian method outlined in Section 2. The results are shown in Figs 2 and 3, respectively. If we characterize the results by the median and 1σ confidence range of the mass estimates, the Davies & Beasor (2018) algorithm finds $M_l/M_\odot = 9.34_{-0.69}^{+0.58}$ and $M_h/M_\odot = 20.01_{-1.84}^{+2.61}$. The results are biased to be higher than the input masses, which is opposite to the sense expected by Davies & Beasor (2018). Our Bayesian approach yields $M_l/M_\odot = 8.44_{-0.65}^{+0.70}$ and $M_h/M_\odot = 18.55_{-1.90}^{+1.95}$, so it is also biased to higher masses, but by a smaller amount. The scatters in the results for M_l are comparable (0.61 versus $0.73 M_\odot$), but the Bayesian estimates of M_h show significantly less scatter (2.62 versus $2.07 M_\odot$). Using different mass limits produces similar results. For example, if we raise the upper limit to $M_h = 21 M_\odot$, we find $M_h/M_\odot = 22.95_{-1.71}^{+2.12}$ for the Davies & Beasor (2018) method and $21.27_{-1.87}^{+2.68}$ for the Bayesian method.

The mass estimates from the two statistical approaches are strongly correlated, as shown in Fig. 4. Simulated data, which lead to an overestimate of M_h by one method, also produce an overestimate by the other method; but the Bayesian method produces mass estimates systematically closer to the input values. Examining the cases with the highest mass estimates, there is a fairly general pattern. The highest mass model star producing a magnitude measurement has a mass close to M_h . The randomly

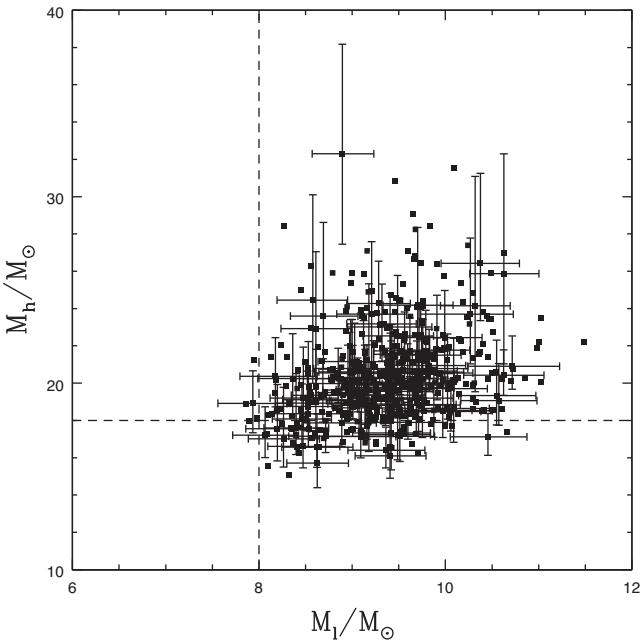


Figure 2. The results for 500 simulated progenitor data sets using the Davies & Beasor (2018) approach to estimating the minimum and maximum progenitor mass. Each case has a point at the median and error bars encompassing 68 per cent ('1 σ ') of the probability are shown for 20 per cent of the trials. The input values are indicated by the dashed lines.

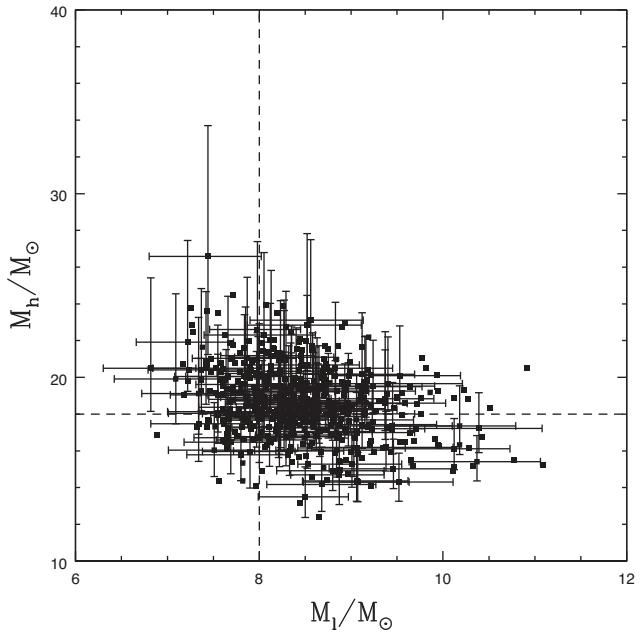


Figure 3. The results for the same 500 simulated progenitor data sets using the Bayesian method presented in Section 2.

selected distance modulus, extinction, bolometric correction, and magnitude error combine to produce a model magnitude that is brighter than the magnitude that would be found using the nominal values for these quantities. Then, when the model is a fit to estimate M_h , the solutions are biased high.

The systems with large mass uncertainties are also the reason why the additional upward correction added by Davies & Beasor (2018)

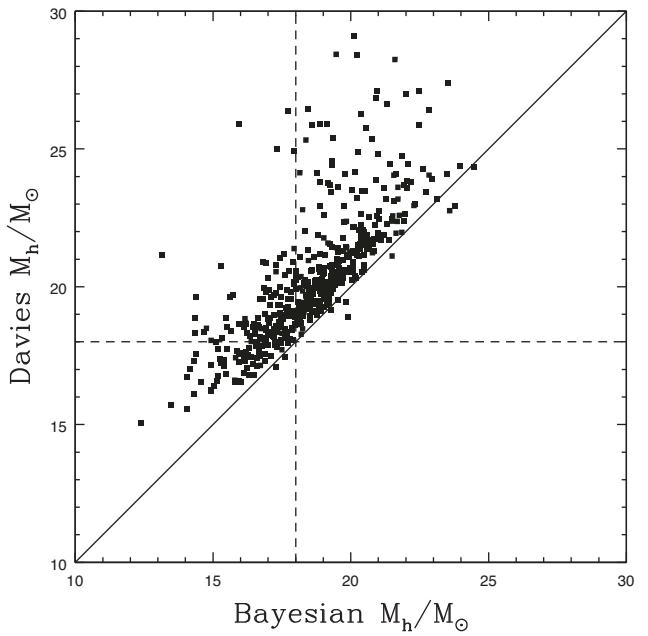


Figure 4. Estimates of M_h for the Bayesian and Davies & Beasor (2018) methods for each of the 500 simulated progenitor data sets. The estimates are strongly correlated, but the Bayesian estimates are systematically lower and closer to the input value. The dashed lines mark the input value of $M_h = 18 M_\odot$ and the diagonal line corresponds to equal mass estimates.

should not be included. If the mass uncertainties are sufficiently small, then the value of M_h estimated from a finite-sized sample will be an underestimate of the true limit as they argue. But this holds only until the typical offset of the highest mass progenitor in the sample from the true upper limit is comparable to the uncertainties in the masses. Once the uncertainties are larger, the analysis is subject to a form of Malmquist bias, where it becomes increasingly likely that an intrinsically lower mass (or equivalently, lower luminosity) star will be interpreted as a star above the true upper mass limit. Based on our Monte Carlo simulations, this appears to be the regime appropriate to the existing progenitor data.

For both approaches, there is a correlation between the uncertainties in M_h and the degree to which M_h is overestimated, as shown in Fig. 5 for the Bayesian simulations. The larger the uncertainties towards larger masses are, the more the mean of the estimator is biased high. For the $M_h = 18 M_\odot$ simulations and the uncertainties estimated for the actual progenitor sample below, these correlations would predict that the Davies & Beasor (2018) estimate is biased high by $(2.6 \pm 1.6) M_\odot$ and the Bayesian estimate is biased high by $(3.3 \pm 0.8) M_\odot$.

Finally, if we analyse the actual progenitor data, we find the results given in Table 1. We include the estimates from Smartt (2015) and Davies & Beasor (2018) using the Eldridge & Tout (2004) mass-luminosity relation (labelled ET04) and then the results using the implementation of both the Davies & Beasor (2018) method (labelled Davies) and the Bayesian method (labelled Bayes). We then repeat the results on the Sukhbold et al. (2018) scale (labelled S18), although these are simply an offset in the mass scale. Including our knowledge that these analyses yield estimates of M_h that are biased to be high, we see that the existence of the red supergiant problem is quite secure unless the maximum mass of stars that undergo core collapse as RSGs can be driven below $20 M_\odot$.

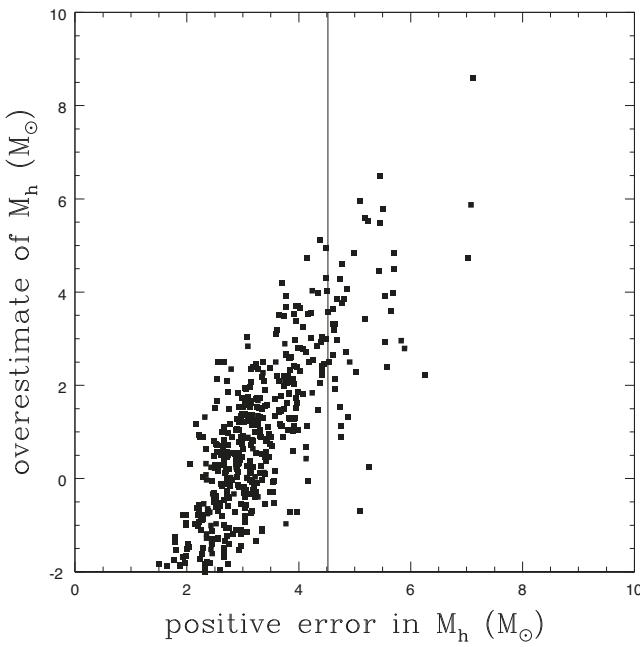


Figure 5. The correlation between the estimate of the positive (upwards) error bar on the Bayesian estimate of M_h and the overestimate of M_h . The larger the estimated uncertainty, the larger the overestimate of M_h . A vertical bar shows the positive uncertainty estimate for the Bayesian analysis of the actual sample. The Davies & Beasor (2018) method shows a similar correlation.

Table 1. Mass limits for the progenitor sample.

Model	$M(L)$	M_h/M_\odot	M_h/M_\odot
Smartt (2015)	ET04	$9.5^{+0.5}_{-2.0}$	$16.5^{+2.5}_{-2.5}$
Davies & Beasor (2018)	ET04	$7.5^{+0.3}_{-0.2}$	$19.0^{+2.5}_{-1.3}$
Davies	ET04	$7.49^{+0.25}_{-0.27}$	$19.05^{+2.22}_{-1.30}$
Bayes	ET04	$6.30^{+0.48}_{-0.54}$	$19.01^{+4.04}_{-2.04}$
Smartt (2015)	S18	$10.0^{+0.5}_{-1.5}$	$18.5^{+3.0}_{-4.0}$
Davies	S18	$8.38^{+0.28}_{-0.30}$	$21.33^{+2.48}_{-1.46}$
Bayes	S18	$7.06^{+0.54}_{-0.61}$	$21.28^{+4.52}_{-2.28}$

Note. The mass–luminosity relation $M(L)$ is either that of ET04 (Eldridge & Tout 2004) or S18 (Sukhbold et al. 2018), where the latter yields a mass about 12% higher for the same luminosity. The uncertainties are 1σ except for the Smartt (2015) estimates which are at 95% confidence.

ACKNOWLEDGEMENTS

We thank Ben Davies for his comments. CSK is supported by NSF grants AST-1908570 and AST-1814440.

REFERENCES

Adams S. M., Kochanek C. S., Gerke J. R., Stanek K. Z., Dai X., 2017a, *MNRAS*, 468, 4968
 Adams S. M., Kochanek C. S., Gerke J. R., Stanek K. Z., 2017b, *MNRAS*, 469, 1445
 Auchettl K., Lopez L. A., Badenes C., Beacom J. F., Ramirez-Ruiz E., Holland-Ashford T., 2019, *ApJ*, 871, 64
 Beasor E. R., Davies B., 2016, *MNRAS*, 463, 1269
 Davies B., Beasor E. R., 2018, *MNRAS*, 474, 2116
 Díaz-Rodríguez M., Murphy J. W., Rubin D. A., Dolphin A. E., Williams B. F., Dalcanton J. J., 2018, *ApJ*, 861, 92
 Dwarkadas V. V., 2014, *MNRAS*, 440, 1917
 Eldridge J. J., Tout C. A., 2004, *MNRAS*, 348, 201
 Gerke J. R., Kochanek C. S., Stanek K. Z., 2015, *MNRAS*, 450, 3289
 Groh J. H., Meynet G., Georgy C., Ekström S., 2013, *A&A*, 558, A131
 Kochanek C. S., 2014, *ApJ*, 785, 28
 Kochanek C. S., 2015, *MNRAS*, 446, 1213
 Kochanek C. S., Beacom J. F., Kistler M. D., Prieto J. L., Stanek K. Z., Thompson T. A., Yüksel H., 2008, *ApJ*, 684, 1336
 Kochanek C. S., Khan R., Dai X., 2012, *ApJ*, 759, 20
 O’Connor E., Ott C. D., 2011, *ApJ*, 730, 70
 Pejcha O., Thompson T. A., 2015, *ApJ*, 801, 90
 Salpeter E. E., 1955, *ApJ*, 121, 161
 Schaller G., Schaerer D., Meynet G., Maeder A., 1992, *A&AS*, 96, 269
 Smartt S. J., 2009, *ARA&A*, 47, 63
 Smartt S. J., 2015, *PASA*, 32
 Smartt S. J., Eldridge J. J., Crockett R. M., Maund J. R., 2009, *MNRAS*, 395, 1409
 Sukhbold T., Ertl T., Woosley S. E., Brown J. M., Janka H.-T., 2016, *ApJ*, 821, 38
 Sukhbold T., Woosley S. E., Heger A., 2018, *ApJ*, 860, 93
 Ugliano M., Janka H.-T., Marek A., Arcones A., 2012, *ApJ*, 757, 69
 Walmswell J. J., Eldridge J. J., 2012, *MNRAS*, 419, 2054
 Zapartas E., de Mink S. E., Justham S., Smith N., Renzo M., de Koter A., 2020, preprint ([arXiv:2002.07230](https://arxiv.org/abs/2002.07230))

This paper has been typeset from a \TeX / \LaTeX file prepared by the author.