

Pulling out all the stops: searching for RPV SUSY with stop-jets

Yang Bai,^a Andrey Katz^b and Brock Tweedie^{c,d}

^a*Department of Physics, University of Wisconsin,
1150 University Avenue, Madison, WI 53706, U.S.A.*

^b*Center for the Fundamental Laws of Nature, Jefferson Physical Laboratory, Harvard University,
17 Oxford Street, Cambridge, MA 02138, U.S.A.*

^c*Physics Department, Boston University,
590 Commonwealth Avenue, Boston, MA 02215, U.S.A.*

^d*PITT PACC, Department of Physics and Astronomy, University of Pittsburgh,
3941 O'Hara Street, Pittsburgh, PA 15260, U.S.A.*

E-mail: yangbai@physics.wisc.edu, andrey@physics.harvard.edu,
btweedie@bu.edu

ABSTRACT: If the lighter stop eigenstate decays directly to two jets via baryonic R-parity violation, it could have escaped existing LHC and Tevatron searches in four-jet events, even for masses as small as 100 GeV. In order to recover sensitivity in the face of increasingly harsh trigger requirements at the LHC, we propose a search for stop pairs in the highly-boosted regime, using the approaches of jet substructure. We demonstrate that the four-jet triggers can be completely bypassed by using inclusive jet- H_T triggers, and that the resulting QCD continuum background can be processed by substructure methods into a featureless spectrum suitable for a data-driven bump-hunt down to 100 GeV. We estimate that the LHC 8 TeV run is sensitive to 100 GeV stops with decays of any flavor at better than 5σ -level, and could place exclusions up to 300 GeV or higher. Assuming Minimal Flavor Violation and running a b -tagged analysis, exclusion reach may extend up to nearly 400 GeV. Longer-term, the 14 TeV LHC at 300 fb^{-1} could extend these mass limits by a factor of two, while continuing to improve sensitivity in the 100 GeV region.

KEYWORDS: Supersymmetry Phenomenology, Jets

ARXIV EPRINT: [1309.6631](https://arxiv.org/abs/1309.6631)

Contents

1	Introduction	1
2	Jet substructure techniques and analysis cuts	4
3	Search strategies and sensitivity estimates	9
4	Conclusions	14
A	Simulation details	15
B	Supplementary results	18

1 Introduction

Well over one hundred null searches for supersymmetric signals at the LHC have now been completed, leaving us to ponder the fate of naturalness after the Higgs boson discovery [1, 2]. As we await the energy upgrade, much attention is being focused on the possibility that supersymmetric particles *have* been produced at the LHC, but for one reason or another are buried amidst the copious Standard Model backgrounds. In this paper, we turn to one of the simplest of these possibilities: the lightest superpartner is a stop eigenstate, and decays promptly to two quarks through a baryonic R-parity violating (RPV) coupling.

Such a situation is in fact quite well-motivated within the context of “effective” or “natural” supersymmetry [3–7], where only the superparticles required to regulate the Higgs mass need to be within the immediate reach of the LHC, and the detailed spectrum and dynamics above the multi-TeV scale can be left unspecified. In this framework, third generation squarks can be amongst the lightest superparticles, forcing us to take seriously the possibility that a stop eigenstate sits at or near the bottom of the SUSY spectrum. R-parity conservation is still often assumed in natural SUSY, but its violation must also be considered. The main motivations for R-parity are stabilization of a supersymmetric dark matter candidate and suppression of proton decay. The dark matter motivation is attractive, but the issues of dark matter and electroweak naturalness might easily be decoupled. The proton decay motivation does not require R-parity per se, but only that baryon-number violation (BNV) and lepton-number violation (LNV) are not simultaneously active. Moreover, the motivation for R-parity is further weakened by the presence of R-parity-conserving operators at dimension-5 with simultaneous BNV and LNV. In a relatively model-agnostic approach with a multi-TeV cutoff, these operators in any case force us to consider proton-stabilizing symmetries other than R-parity. RPV is also sometimes considered dangerous given precision tests, in particular $K - \bar{K}$ and $n - \bar{n}$ mixing in the

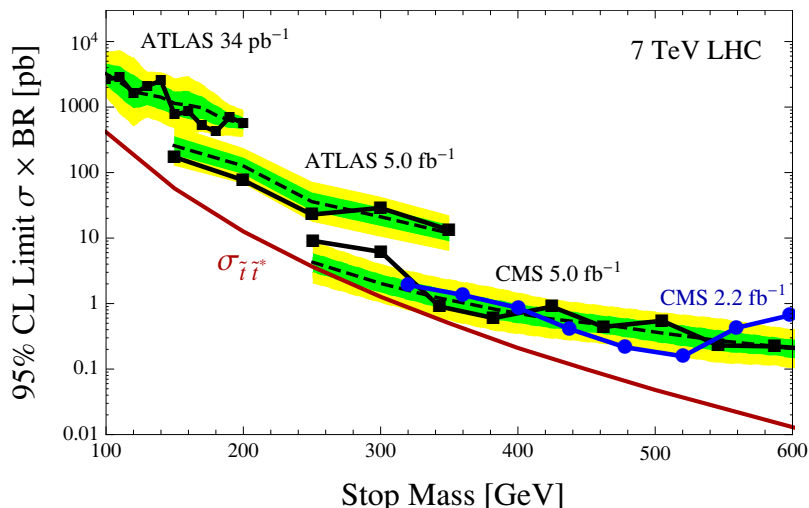


Figure 1. Existing constraints on $pp \rightarrow \tilde{t}\tilde{t}^* \rightarrow 4j$ from the LHC, reinterpreting the results of [10–13] to account for stop acceptances relative to coloron or hyperpion acceptances.

BNV case. However, couplings that are small enough to avoid these constraints are still usually large enough to allow for prompt RPV decays from the perspective of the LHC [8] (see [9] for a complete list of constraints).

The idea that SUSY might be both “natural” and violate R-parity has received increasing attention in the past few years [6, 14–25], exactly because the relatively small production cross sections and non-canonical decay topologies lend themselves to evading LHC searches. Many spectra and choices of active RPV operators exist, and the nontrivial flavor structures of the couplings must be specified. In some cases, especially those involving leptonic RPV, exclusion or discovery prospects are quite good (see [14] for proposals and recast limits, and e.g. [26, 27] for some recent experimental searches). Successful baryonic RPV searches have also been performed, such as LSP gluinos decaying fully hadronically into multiple jets [28–31]. However, an LSP stop decaying promptly to two jets via baryonic RPV has proven particularly evasive. Stop pairs are produced at much smaller rates than gluino pairs of comparable mass, and the lower-multiplicity 4-jet final state is even more difficult to disentangle from the pure QCD backgrounds. Another major complicating aspect at the LHC is the multijet triggers, which can heavily prescale-away the signatures of stops lighter than several hundred GeV. Some of the best current direct limits actually come from LEP, which rules out $m_{\tilde{t}} \lesssim 90$ GeV [32]. A recent search at the Tevatron extends this limit up to only about 100 GeV [33]. However, so far, direct searches for pair-production of dijet resonances at the LHC have failed to reach the sensitivity necessary to place constraints for any stop mass [10–13]. A snapshot of the current situation can be seen in figure 1. In fact, the inevitable rise of trigger thresholds with instantaneous luminosity and beam energy leaves us to wonder whether the LHC will *ever* be sensitive to this signal. At the very least, this trend suggests that masses near the current limit of 100 GeV might be left unexplored.¹

¹For recent projections for the long-term LHC, which begin to achieve exclusion reach but nonetheless do not pursue signals below 300 GeV, see the recent Snowmass study [34].

One way around these difficulties is to search for the stop as a dijet resonance produced in the decays of heavier colored superparticles, such as gluinos [35] or sbottoms [8] (or possibly the heavier stop eigenstate), or to simply set bounds using the associated leptonic activity and high H_T of these decays [36–39]. Naturalness suggests that these colored superparticles should also not be far above 1 TeV, and might be produced with observable rates. It is also possible to invoke Minimal Flavor Violation (MFV), which suggests that stops dominantly decay (with a branching ratio $\simeq 95\%$) into $\bar{b}\bar{s}$ or $\bar{b}\bar{d}$ [15]. It was pointed out in [40] that incorporating b -tagging into the triggering might allow the direct stop pair signal to write to tape with higher efficiency, and subsequent kinematic analysis can discriminate it from flavored backgrounds. The use of b -tagging is also usually considered to help identify stop production in gluino decay. However, at present, there seems to be no strategy at the LHC that guarantees full coverage of the general $\tilde{t}\tilde{t}^* \rightarrow 4j$ signal, exploiting neither spectrum-dependent nor flavor-dependent features of the theory.

Rather than consider the generic signal lost, we here exploit what is now a well-worn trick [41–43]: we focus on “boosted” production at high transverse momentum, and apply the methods of jet substructure to “stop-jets,” each containing an entire stop decay. Indeed, analyses of this type have already been applied in the search for light RPV gluinos [29, 44, 45], and have also been studied for color-octet scalars [46], RPV neutralinos from squark decays [47], and the aforementioned boosted stops from gluino decays [35]. The advantages of a boosted/substructure analysis are manyfold, including better S/B , automatic resolution of combinatoric ambiguities, and improved mass resolution due to more complete decay radiation containment and rejection of uncorrelated soft radiation. But probably the primary advantage for a stop pair search is in how a substructure treatment can process the continuum QCD background. The standard resolved 4-jet analyses partition the leading jets into two pairs so as to form two stop candidates, either minimizing the candidate mass asymmetry (CMS) or minimizing a ΔR measure for the two decays (ATLAS). The stop is searched for as a bump in the average pair-mass spectrum. In either case, the requirement of four well-separated and high- p_T jets passing the triggers, as well as other downstream analysis cuts, leads to a highly-shaped background QCD spectrum with a pronounced “trigger turn-on” peaking at 100–200 GeV. As a consequence, the most recent CMS [11] (ATLAS [13]) analysis of 2011 data does not even search for masses below 250 GeV (150 GeV).

By contrast, as we will see, a carefully constructed jet substructure approach potentially leads to a largely featureless average-mass spectrum for the QCD background from a few 10’s of GeV up. This is because jet substructure, unlike traditional jet analyses, gives us far more flexibility in assigning sprays of hadrons to individual hard “quarks” or “gluons.” In particular, we can make these assignments in a way much closer to QCD itself, which is approximately a scale-invariant theory. To best exploit this feature, we will also capitalize on the summed jet- H_T trigger rather than triggers that count specific numbers of jets above some threshold. The jet- H_T trigger, which exists in some form in both CMS and ATLAS, probably has the least sensitivity to precisely how the event activity groups into standard LHC jets.

With an unbiased background spectrum, a bump-hunt becomes much more tractable, regardless of what stop mass we consider. Either a parametrized spectral fit or kinematic

sideband approaches like the ABCD method can be employed to determine the QCD background, and we explore the viability of several such approaches, including some novel ones. We find that 100 GeV stops should be visible in the 2012 data set with better than 5σ statistical significance. Exclusion reach should extend up to more than 300 GeV. We also consider what might be possible in the MFV case, using a b -tagged version of the analysis. S/B improves dramatically, for example yielding 10σ -level statistical sensitivity to 100 GeV stops, and exclusion up to almost 400 GeV. For the longer-term Run II of the LHC, with a projected 300 fb^{-1} at 14 TeV, we estimate untagged exclusions extending up to 650 GeV, with discovery-level sensitivity or better between 100 GeV and 500 GeV.

Our paper is organized as follows. In the next section, we discuss our substructure procedures, and explore the response of the signal and QCD continuum in simulation data. We present our data analysis techniques and final sensitivity estimates in section 3. We conclude in section 4. Two appendices contain more details of our simulations and more in-depth substructure studies.

2 Jet substructure techniques and analysis cuts

The fundamental obstacle to any multijet search at the LHC is the overwhelming production rate for continuum QCD backgrounds. The dynamics of QCD is approximately scale-invariant at high energy. Convolved with smoothly-falling parton distribution functions, a good stage is set for extracting signals with sharply-localized mass features, such as RPV stop pairs. However, the standard approaches to reconstructing QCD events are far from scale-invariant. “Jets” as usually defined are tied to a dimensionful p_T threshold, as well as a dimensionless jet radius R . For example, in a multijet search with $p_T(j) > 100 \text{ GeV}$ and $R = 0.5$, there is an absolute minimum invariant mass between jet pairs of approximately 50 GeV. Any QCD background invariant mass spectrum constructed from such jets will start at 50 GeV, rise at higher invariant masses as the efficiency turns on, and only then turn over into a smoothly-falling shape. Searches for low- S/B features in the broad turn-on region or at the peak are often not considered, as the precise signal and background shapes in these regions have high sensitivity to reconstruction uncertainties.

Indeed, these kinds of considerations have limited the range of applicability of current LHC stop searches based on the multijet approach, as seen in figure 1. To control triggering rates as the LHC continues to increase instantaneous luminosity and energy, the jet p_T thresholds are gradually increasing. For example, to keep 4-jet triggering rates at their 2012 levels during the projected 300 fb^{-1} run at 14 TeV, offline $p_T(j)$ thresholds would need to be increased to 150–200 GeV. Stops sitting near the current exclusion limit of $m = 100 \text{ GeV}$ might therefore never be visible. Even with 2011 data, only the very low-luminosity ATLAS search [12] was capable of probing masses near that limit, but was not sensitive to stop pair cross sections. Despite the fact that the cross section is becoming quite large near 100 GeV, basic triggering and analysis cuts at nominal LHC luminosities are carving away the signal.

As has been pointed out in [48, 49], reconstruction-induced biases can be highly ameliorated by applying modern approaches to jet substructure on “fat-jets.” Jet substructure can be less tied to fixed p_T or ΔR thresholds for the individual reconstructed hard par-

tons, and can put more focus on dimensionless quantities such as ratios of masses or p_T 's. The “price” for applying a substructure-based search is that we must work in the boosted region of production phase space, which might still represent only a small fraction of the total signal. Nonetheless, as we will see, the tradeoff is more than worthwhile. In any case, existing 4-jet searches are already capitalizing on the boosted or semi-boosted region of phase space, which offers the additional advantages of better S/B and much-reduced combinatoric ambiguities.

We are in the midst of an ongoing boom of jet substructure ideas, many of which can have applicability to the boosted stop pair signal. We here focus on the relative- p_T declustering method used in the JHU top-tagger [50] and the diboson-jet tagger of [51], which is itself based on the “BDRS” method of [43]. These methods identify localized clusters of energy within the fat-jets as “subjets,” and subsequently treat these objects similar to ordinary QCD jets. However, before proceeding, we emphasize that ideas such as (but not limited to) pruning [52], N-subjettiness [53], template overlaps [54], shower deconstruction [55], and Q-jets [56] could all worth more dedicated study in this context, and might lead to further improvements.

For our fat-jet clustering radius, we choose a very large value of $R = 1.5$, roughly giving us hemisphere-sized jets. The motivation for such a large radius is twofold. First, the large catchment area gives us a broader reach in stop masses. For given momentum thresholds, heavier stops will be able to pass with less boost, and therefore will have more widely-separated decay products. Heavier stops also have much smaller cross section, and the very highly-boosted portion of phase space may be too poorly-populated to be exploited. In fact, for much of the mass range probed in this paper, the quarks in the stop decays are well-separated enough to be reconstructed as individual jets of normal radius. However, we find it interesting that these masses can alternatively be covered using substructure procedures, raising the question of whether a traditional analysis is even necessary. The second motivation for the large radius is more subtle, and has to do with how the continuum QCD background is processed. Interplaying with momentum and substructure cuts, the jet radius sets a maximum mass (analogous to the *minimum* mass for normal jet-pair masses described above). Larger radii push the turn-off of the mass spectrum out to higher values, creating a less steeply-falling background. Larger radii also somewhat improve S/B . This is because enlarging the jet radius increases sensitivity to wider-angle radiation, and provides a primitive form of color discrimination. At high momenta, stop pair production is dominated by $q\bar{q} \rightarrow \tilde{t}\tilde{t}^*$, whereas the QCD background contains more highly-colored processes such as $qg \rightarrow qggg$. (For a comparison of different R 's, see appendix B.)

Fixing $R = 1.5$, we cluster jets in each event using the Cambridge/Aachen (C/A) algorithm [57, 58] as implemented in `FastJet3` [59], and study the leading two jets within $|\eta| < 2.5$. Then we iteratively undo the clustering stages individually within each jet. Now viewed in reverse as a splitting, at each stage we can look at the two branches “ a ” and “ b ” and decide whether the splitting should be considered “hard” or “soft” according to some prescription. For example, a soft splitting might resolve a low- p_T collection of radiation at the jet’s periphery. In our specific implementation, derived directly from the first declustering stage of the diboson-jet-tagger of ref. [51], a splitting is considered hard if both a and b carry

appreciable p_T relative to the original fat-jet, and if the m/p_T ratio of the softer branch is not too large. Explicitly, assuming $p_T(b) < p_T(a)$, we require $p_T(b)/p_T(\text{fat-jet}) > 0.1$ and $m(b)/p_T(b) < 0.3$. Otherwise the splitting is considered soft, in which case the lower- p_T branch is thrown away, and the declustering is continued along the surviving branch. The procedure is repeated until a hard splitting is encountered, or no more jet constituents remain (in which case the entire event is vetoed). The two branches at the hard splitting are our two subjets, representing our assignment of the jet's radiation to the two quarks in the stop decay.² The reconstructed stop candidate is the four-vector-sum of these two subjets.

Most of our ability to discriminate stop pair production from ordinary QCD stems from the fact that the stop events contain two subjet-pair resonances of equal mass. Therefore, instead of looking for a bump within the distribution of the individual stop candidate masses m_1 and m_2 , we look for a bump in the joint distribution of (m_1, m_2) . Practically, we can turn this into a 1D bump-hunt by first focusing on the region of small mass asymmetry $A \equiv |m_1 - m_2|/(m_1 + m_2)$, and then constructing the spectrum of the averaged mass $m_{\text{avg}} \equiv (m_1 + m_2)/2$. We pick a nominal mass asymmetry threshold of 10%.

We have found that S/B can be further purified by a handful of additional cuts. In particular, we can place a cut on the stop-pair CM-frame production angle (as is done in [12, 13]), formed by actively boosting the entire 4-subjet system to rest in the lab frame and measuring the angle of either stop candidate with respect to the beamline. We call this angle θ^* , and place a cut $|\cos \theta^*| < 0.3$. Finally, we exploit the fact that energy splittings within QCD jets tend to be very asymmetric, whereas in stop decays they tend to be more democratic. Within each stop pair, we demand that the p_T 's of subjets a and b relative to one another satisfy $\min[p_T(a), p_T(b)]/\max[p_T(a), p_T(b)] > 0.3$.

The final ingredient needed to define an analysis is to establish a trigger for the events. We choose the total jet- H_T trigger of CMS, which sums up the p_T 's of all ordinary jets.³ In ref. [60], based on 2011 data, the H_T trigger was used for a classic jets+ \cancel{E}_T style SUSY search. The trigger was found to be fully efficient for final reconstructed events with $H_T > 750$ GeV, summing over $R = 0.5$ anti- k_T jets with $p_T > 50$ GeV. Anticipating slightly harsher triggers in 2012, we conservatively set our threshold at 900 GeV.⁴

We test our substructure methods and cuts on signal and QCD background simulation samples for the 8 TeV LHC. Full details of their generation can be found in appendix A. We note here that the signal samples are matched up to one jet emission at the production stage, which gives a more accurate H_T spectrum. The QCD continuum simulations are matched up to four hard partons (including b -quarks), using the CKKW-L [61] prescription implemented in Pythia8 [62]. We take a quark-flavor-conscious Durham k_T distance as our merging measure, and a merging threshold of 50 GeV. Other background samples,

²The original BDRS method uses a very similar procedure [43], relying instead on a mass-drop criterion and a somewhat different momentum-asymmetry criterion. While BDRS can be adapted for use in the boosted stop search, there are important caveats which we discuss in more detail in appendix B.

³ATLAS also has a dedicated single-fat-jet trigger, though this is defined with $R = 1.0$. Nonetheless, this trigger could still be used as the basis of a substructure analysis very similar to the one that we explore here.

⁴A looser offline H_T cut may in fact be feasible with 2012 data, and would only improve our results. We thank Keith Ulmer for bringing this point to our attention.

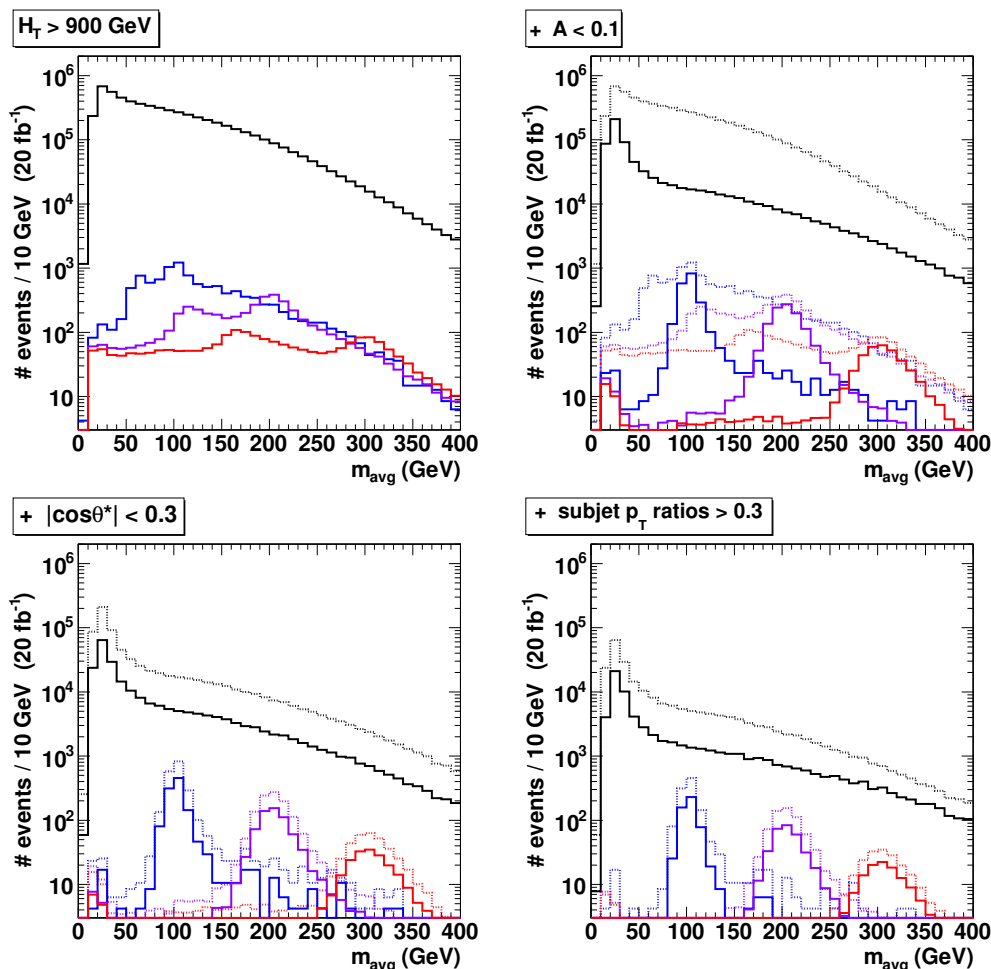


Figure 2. The effects of our cuts on the spectrum of $m_{\text{avg}} \equiv (m_1 + m_2)/2$, defined on declustered fat-jets. From left-to-right, top-to-bottom, cuts are added sequentially. The effect of the preceding cut is shown with dotted histograms for comparison. Background is matched QCD (black), and example stop models are 100 GeV (blue), 200 GeV (purple), and 300 GeV (red).

such as $t\bar{t}$, are generated in MadGraph5 interfaced with Pythia8, or self-contained within Pythia8. All samples are processed through a simple, perfect 0.1×0.1 grid “calorimeter” in η - ϕ space. To prove the robustness of our methods in a high pileup environment, we also introduce pileup events into the simulations and then apply a form of event-wide trimming [63] tailored to pileup removal before the fat-jet clustering stage. (This procedure also heavily reduces the impact of the underlying event and soft ISR.) After declustering the fat-jets, individual subjects are energy-smearred similar to normal LHC jets (see appendix A for further details).

Figure 2 shows the effect of the cut-flow on the m_{avg} spectra of continuum QCD and some example signal mass points. There are several features that are worth noting. First, the ubiquitous turn-on peak is present, but has been pushed down to the $O(10 \text{ GeV})$ scale, far away from our signals. This can easily be understood from the interplay between our




Figure 3. Final spectra of m_{avg} after all cuts, for an untagged analysis (left) and a b -tagged analysis assuming $\text{BR}(\tilde{t} \rightarrow \bar{b}\tilde{d}/\tilde{b}\bar{s}) \simeq 100\%$ and tagging/mistagging rates as described in the text (right). Displayed backgrounds include matched QCD (black), $t\bar{t}$ (pink), and W/Z +jets (green). The matched QCD histogram has been smoothed from the Monte Carlo data, as described in the next section. Displayed example stop models, stacked onto the QCD background, include 100 GeV (blue), 200 GeV (purple), and 300 GeV (red). The lower panels show the S/B ratio relative to QCD, and the bin-by-bin fractional statistical errors on the QCD background expected for the 2012 LHC run. (Note the changes in vertical axes between untagged and tagged.)

H_T cut, the 10% relative- p_T requirement in the declustering, and the size of our calorimeter cells. (Note that the lowest- p_T subjects that we work with are roughly 40 GeV.) For the signal, before the $A < 0.1$ cut, it is possible to see three distinct mass features: the turn-on, the true stop-mass peak, and an intermediate peak near half of the stop mass. The last feature arises from events where one stop is correctly reconstructed, but one of the quarks was lost for the other stop. After the A cut, the stop peaks become very clear. The additional cuts both reduce the overall size of the background and further tighten the signal peaks. Throughout the entire set of cuts, the QCD background in the vicinity of the signals stays “featureless.” The final signal efficiencies relative to the inclusive pair production rates are 5×10^{-5} for 100 GeV and 4×10^{-3} for 300 GeV. For comparison, stop pair acceptances for the standard 4-jet searches are usually at the 10^{-3} – 10^{-2} level.

We show how the signal peaks appear on top of the continuum QCD background in linear scale in the left panel of figure 3, and give an indication of S/B relative to the QCD background’s bin-by-bin statistical errors. It is clear already from this simple plot that the 100 GeV stop should be visible with high significance, and that exclusion reach should extend beyond 200 GeV. We also show on this plot the largest subdominant background contributions, namely $t\bar{t}$ and W/Z +jets. (Other backgrounds such as diboson and single-top can be shown to be far smaller.) We will not directly address the impact of these backgrounds on a search, but they are clearly important. This is especially true for $t\bar{t}$, which contributes a broad bump peaked near 175 GeV, and at a size only $O(1)$ smaller

than our stop signal.⁵ However, the multibody structure of this background is under much better theoretical control than pure QCD, and its normalization could be extracted in the highly orthogonal semileptonic channel. We therefore anticipate that it could be systematically subtracted or accounted for in a constrained fit. Indeed, it can even serve as a useful calibration peak. If it is necessary to further suppress $t\bar{t}$, it might be possible to do so with supplementary substructure cuts that can pick out and reject 3-body features, without highly resculpting the continuum QCD. (E.g., N-subjettiness [53] observables or the dimensionless variables of the HEPTopTagger [64] would be appropriate to study.) Regardless, some degradation of sensitivity in the vicinity of m_t should be expected in reality.

If the RPV coupling obeys MFV, then almost every stop decay will contain a b -quark. It therefore becomes possible to exploit a b -tagged analysis. We show in the right panel of figure 3 the m_{avg} spectra after demanding that at least one of the four subjects is tagged, assuming flat ($b, c, q/g$) tag rates of (60%, 10%, 2%). The S/B (and S/\sqrt{B}) improves dramatically, as does the relative contribution of $t\bar{t}$ to the background budget. Exclusion reach to above 300 GeV already appears highly likely.

These distributions set the stage for our data analysis in the next section.

3 Search strategies and sensitivity estimates

Extraction of a bump on top of a smoothly-falling background spectrum is a classic problem that has appeared many times already at the LHC. Probably the most famous recent application is the observation of the Higgs resonance in the continuum diphoton spectrum [65, 66], but this strategy has also been applied by CMS in its paired dijet resonance searches [10, 11]. Relying solely on the assumption that the background is “featureless,” the observed spectrum can be fit to a parametrized function with or without the addition of a signal bump. The parametrization used by CMS, which we also use here, is a four-parameter function of the form [11]

$$\frac{dP}{dm_{\text{avg}}} = p_0 \times \frac{(1 - m_{\text{avg}}/\sqrt{s})^{p_1}}{(m_{\text{avg}}/\sqrt{s})^{p_2+p_3} \log(m_{\text{avg}}/\sqrt{s})}, \quad (3.1)$$

where p_0, p_1, p_2, p_3 are free parameters and \sqrt{s} is the proton-proton center-of-mass energy (8 TeV for 2012).⁶

There are also many other ways to directly estimate the QCD m_{avg} spectrum from the data, using control regions. The common ABCD method was used by ATLAS in its own paired dijet resonance searches [12, 13]. This method requires defining sideband cuts in two variables. The nominal cuts define region A, and the other three choices

⁵The fact that $t\bar{t}$ is not a larger contribution is perhaps somewhat surprising, given that for $m_{\bar{t}} \simeq m_t$, the inclusive $t\bar{t}$ cross section is about six times larger than $\tilde{t}\tilde{t}^*$. About half of this factor comes from the $t\bar{t}$ all-hadronic branching fraction, since only all-hadronic events are efficient at passing the H_T cut and subsequent substructure cuts. It is also important to realize that for high- p_T central production, the difference in cross sections is not as big. (Asymptotically, the factor of six reduces to a factor of two.) Finally, the large fraction of partial reconstructions with two-body substructure significantly broadens the top peak shape.

⁶It may be possible to develop parametrizations that more directly incorporate the analytic structure of QCD (see [67, 68]), though we have not explored this possibility.

of 2D cuts (nominal+sideband, sideband+nominal, and sideband+sideband) define three signal-depleted regions B, C, and D. Assuming small correlations between the two variables (an assumption that must be justified in Monte Carlo studies), the region-A background spectrum can be derived by taking the bin-by-bin ratios of counts $B \times C/D$. For its ABCD-based search, ATLAS uses the variables A and $|\cos\theta^*|$ to define its four regions. We run our own version of this search, taking sideband cuts $A = [0.1, 0.4]$ and $|\cos\theta^*| = [0.3, 0.8]$.

In addition to the shape-fit and ABCD methods, we also explore two supplemental techniques which may further improve statistical power.

In the first of these, we run a simultaneous shape fit with eq. (3.1) over our nominal signal region and a nearby asymmetry sideband region $A = [0.1, 0.2]$. This increases the data statistics available to the fit by $O(1)$, as well as folding in more discriminating characteristics between the stop signal and QCD. The method is based on the assumption that the QCD spectrum changes very slowly as a function of A when $A \ll 0$, whereas the signal is strongly peaked in A for $m_{\text{avg}} \simeq m_{\bar{t}}$. We observe exactly this behavior in our simulations. To implement the simultaneous fit, we set the QCD fit parameters identical between the two A regions. We can then add into the fit the signal shapes and relative normalizations appropriate to each individual region. This simple two-region fit might also be improved by more finely subdividing in A , and thereby more fully exploiting the nontrivial shape of the signal over this variable. Equivalently, this can be viewed as a full 2D fit over the small- A region of the (m_1, m_2) plane.

The final method that we study (inspired in part by [69]) is based on the assumption that the two fat-jet masses can be considered approximately uncorrelated in background events. If this assumption holds, then the m_{avg} spectrum in the $A \rightarrow 0$ limit can be directly predicted from the spectra of individual fat-jets.⁷ Given a single-jet probability distribution dP/dm_{jet} (where m_{jet} represents the subjet-pair mass), we get

$$\left. \frac{dP}{dm_{\text{avg}}} \right|_{A < A_{\text{cut}} \ll 1} \sim \left[\frac{dP}{dm_{\text{jet}}}(m_{\text{avg}}) \right]^2 \times m_{\text{avg}}. \quad (3.2)$$

The major advantage of this method is that the statistics available for measuring dP/dm_{jet} are enormous, since no tight cuts should be placed on the event-by-event jet-mass asymmetry. However, the uncorrelated assumption needs to be carefully studied in simulation. (It may also be tested to some extent in data, by inverting some of our final cuts.) To keep the kinematics similar to our final signal region, we measure the dP/dm_{jet} spectrum in a control region with only the H_T and $|\cos\theta^*|$ cuts in place. We pick a random jet amongst the leading two, and further demand the $\min[p_T(a), p_T(b)]/\max[p_T(a), p_T(b)]$ cut

⁷Ignoring correlations in emissions that cross between the two jets, at leading-order the main source of possible correlations comes from the $qq/qg/gg$ content of the initial state. Given our H_T cut, this is in fact an approximately 1/2/1 admixture, allowing us to meaningfully define an “average jet” by averaging the mass distributions of q and g jets. The construction that we present is only very mildly sensitive to deviations from this special admixture, due to a combination of two favorable features: a non-factorizable gg component of the $qq/qg/gg$ admixture would itself look like an added contribution to the predicted m_{avg} spectrum of equal parts q and g , and at large m/p_T the q -jet and g -jet mass distributions are actually fairly similar up to an overall normalization. The non-factorizable m_{avg} shape corrections are only quadratically sensitive to the fractional differences between q -jet and g -jet shapes.




Figure 4. Relative agreement between different methods of estimating the continuum QCD m_{avg} spectrum, for untagged (left) and b -tagged (right) analyses. Here, “1” is defined as the prediction from our Monte Carlo data, smoothed with the exponential of a fifth-degree polynomial. The black histogram is the original Monte Carlo data, with error bars estimated from the quadrature-sum of event weights. The four curves correspond to four estimation methods: shape fit applied to the original MC data (green), ABCD (purple), asymmetry-sideband (blue), and single-jet template (red).

for only that jet. (The signal contamination in this single-jet control sample is negligibly small.) Once the dP/dm_{avg} spectrum estimate is derived, we use it as a template for a one-parameter normalization fit to our signal-region spectrum, with or without a signal bump added. In practice, we approximate dP/dm_{jet} with a very finely-binned spectrum measured from our “data,” apply the above transformation, and then integrate back to our nominal m_{avg} binning (10 GeV).

Since our four search methods (shape fit, ABCD, asymmetry-sideband, and single-jet template) all rely on assumptions about the behavior of the QCD background, we can first get some sense of the validity of those assumptions. To do so, we compare the signal-region QCD spectrum to the predictions of the four methods, without signal. Here and below, we focus on the mass range $m_{\text{avg}} = [60, 500]$ GeV. Also, due to the finite statistics of our QCD Monte Carlo sample, which is actually comparable to a 2012-like data set, we smooth out the spectra in the signal and control regions by fitting to the exponential of a fifth-degree polynomial. These fits generally have high goodness-of-fit probability.⁸ (We do not apply smoothing to the single-jet spectrum, which has much higher statistics.) We show the results of the comparison in figure 4, including the original Monte Carlo data with its statistical errors. For $m_{\text{avg}} \lesssim 300$ GeV, the agreement is generally better than 5% (10% with the smaller-statistics b -tagged sample). At higher masses, more pronounced disagreements

⁸A similar parametrization could also be considered for the shape fit search method, or alternatively we could have smoothed our Monte Carlo data using the functional form of eq. (3.1). In fact, we have found that the latter does indeed furnish a high-probability fit. However, we have explicitly chosen different functions for smoothing the raw Monte Carlo and for fitting the derived pseudodata, in order to help prevent spuriously good pseudodata fits.

develop, but their significance is likely not large given the growing error bars. This is also entering the m_{avg} region where the subjects are nearly separated by $\Delta R = 1.5$ and the spectrum is turning off, in which case we might not be surprised to find more sensitivity to the different control region cuts used in the estimation procedures. Needless to say, a higher-statistics simulation would be useful here to better gauge the level of agreement at both higher and lower masses. Still, given the encouraging agreement over the mass range in which we will shortly find our best sensitivity, we can proceed with our analyses without fear that their underlying assumptions are grossly invalid, at least for matched QCD.⁹

We now apply the four data-driven methods to search for the stop signal in our simulations, using the matched QCD background. For all four methods, we search for a signal using χ^2 differences as a discriminator. In the case of the fits, we take the χ^2 difference between background-only and signal+background fits, using \sqrt{N} error bars for the observed bin counts. For the ABCD method, which directly predicts the background spectrum with no free parameters, we construct χ^2 bin-by-bin by combining in quadrature the statistical errors of the observed spectrum and the predicted spectrum. (The latter is itself derived by simple propagation-of-errors from the B, C, and D bin counts.) In all cases, the signal strength and shape are fixed, and systematic errors are not assessed. Our results should therefore be indicative of what can be accomplished in the limit of small systematics.¹⁰

We use our (smoothed) background simulations and signal+background simulations as the basis of a large number of pseudoexperiments corresponding to 20 fb^{-1} at 8 TeV. We run our various search strategies on these pseudoexperiments to build up $\Delta\chi^2$ distributions. We find these distributions to be highly Gaussian, at least out to about three standard deviations. To better parametrize their separation, we therefore fit the $\Delta\chi^2$ distributions to Gaussians. From these we can derive a “median discovery significance” and a “median exclusion significance.” These are the distance between the medians of the two $\Delta\chi^2$ distributions, as measured in units of the background-only pseudoexperiments’ σ (for discovery) or the signal+background pseudoexperiments’ σ (for exclusion). We can set a benchmark for discovery at $5\sigma_B$ separation, not accounting for the look-elsewhere effect. We can set a benchmark for exclusion at $2\sigma_{S+B}$, which is practically equivalent to the usual 95% CL_S criterion. For exclusion, we also consider fluctuations at $\pm 1\sigma_B$ about the background-only median, and recompute the exclusion level in σ_{S+B} units.

Figures 5 and 6 show the final results for untagged and b -tagged analyses, respectively. It is clear that all four search methods perform comparably, but that the single-jet tem-

⁹We have also rerun the untagged comparison on a larger background sample generated wholly within `Pythia8`, based on showered $2 \rightarrow 2$ production without matching. The shape fit and ABCD methods continue to improve across the full mass range, generally agreeing with the signal region spectrum to better than 2% for $m_{avg} < 350 \text{ GeV}$ (and still consistent within MC errors). The asymmetry-sideband agreement exhibits a -20% dip above 300 GeV, similar to the matched case, but with much higher significance. Interestingly, the single-jet template also develops a broad +20% discrepancy above 300 GeV, which appears to be more severe than for the matched case. Again, these features occur in regions where our sensitivity is in any case becoming poor, but they would need to be more carefully investigated in a fully realistic search.

¹⁰Systematic errors affecting the signal shape and normalization will include the jet/subject energy scales and resolution uncertainties, and probably PDF and scale errors for the high- p_T cross section. Possible systematics affecting our different data-driven estimates of the QCD spectrum would need to be investigated with higher-resolution Monte Carlo simulations with a full detector mock-up and/or in control data.

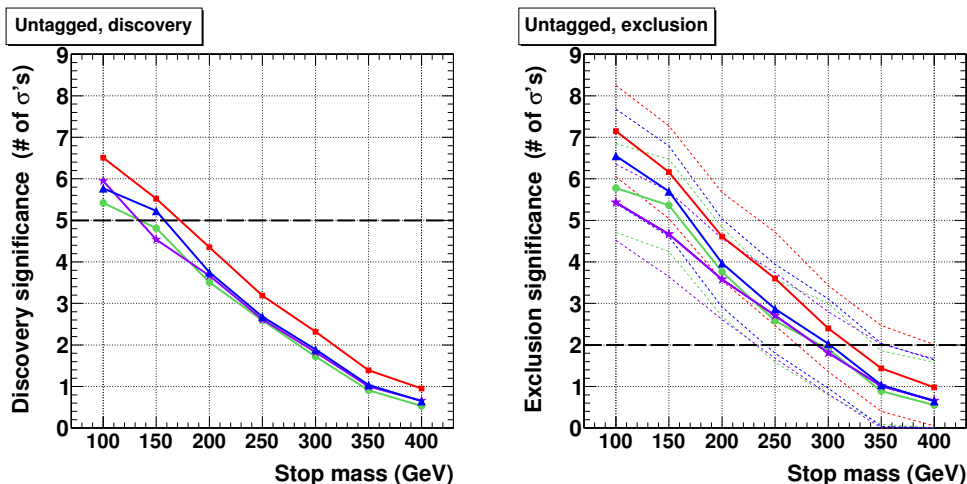


Figure 5. Results of our untagged search estimates for 20 fb^{-1} at LHC8, displaying median discovery significance (left) and exclusion significance (right). The four curves correspond to four data-driven QCD background estimation methods: shape fit (green), ABCD (purple), asymmetry-sideband (blue), and single-jet template (red). In the exclusion significance plot, we also indicate the $\pm 1\sigma$ variation expected due to background statistical fluctuations.




Figure 6. Results of our b -tagged search estimates for 20 fb^{-1} at LHC8, displaying median discovery significance (left) and exclusion significance (right). The four curves correspond to four data-driven QCD background estimation methods: shape fit (green), ABCD (purple), asymmetry-sideband (blue), and single-jet template (red). In the exclusion significance plot, we also indicate the $\pm 1\sigma$ variation expected due to background statistical fluctuations.

plate method tends to edge out the other three, and that the asymmetry-sideband method offers a small but consistent improvement over the simple shape fit. (In fact, for exclusion significance, the single-jet template method gives results very close to what would be inferred with a naive S/\sqrt{B} analysis with optimized mass windows.) The similarity of the results is encouraging, and suggests that experimentalists will have many alternative choices for performing cross-checks of a tentative signal, or as fall-back options if any of

these data-driven methods turns out to be unreliable. From figure 5, which shows the untagged analysis, we see that stops less than about 175 GeV could be discovered, and stops less than about 320 GeV could be excluded. For the b -tagged analysis in figure 6, masses below 250 GeV are discoverable, and exclusion sensitivity extends to nearly 400 GeV. We note that this analysis was run without re-optimization of our cuts, so it might be possible to construct an even more sensitive search. It may also be possible to make even further gains by considering a double- b -tagged search.

Looking ahead, we have also run versions of these analyses on 14 TeV simulations, assuming 300 fb^{-1} luminosity, and for simplicity neglecting pileup. Here, we have used a summed-jet H_T cut of 1600 GeV, which keeps the rate approximately the same as the 900 GeV threshold under 2012 conditions (assuming quadrupled instantaneous luminosity at 14 TeV). The 100 GeV untagged signal remains visible, with statistical significance slightly better than our 2012 estimate, though with approximately 2–3 times smaller S/B . The discoverable range expands up to about 500 GeV, and masses of 200–300 GeV would be visible at the 10σ -level. Exclusion should extend up to 650 GeV. This last finding is comparable to that of the recent Snowmass 2013 report [34], which uses traditional jet reconstruction methods and a highly approximate background estimate. However, that search assumes 2012-like jet p_T cuts, and even then is limited to the mass range above 300 GeV. By contrast, in our jet substructure version of the search there is practically no low-mass cutoff on the search range, with masses from 100 GeV to $O(\text{TeV})$ covered by a single analysis strategy.

4 Conclusions

In this paper, we have addressed what has been believed to be one of the most difficult supersymmetry signatures at hadron colliders, and demonstrated that it may nonetheless be made highly visible using the tools of jet substructure. Besides serving as a crucial supplement to the LHC’s broad-based program for testing naturalness, this result, if reproducible in a realistic analysis on actual LHC data, will serve as a benchmark for fully jetty searches. The implications extend well beyond just RPV supersymmetry. Thus far, multijet searches at the LHC have successfully placed constraints on the pair-production a variety of colored objects, whose prompt decays contain neither leptons, neutrinos, nor other invisible particles [10–13, 28–31]. However, compared to the LSP stop with baryonic RPV, which is a color-triplet scalar undergoing a two-body decay, these searches have always relied on more color, more spin, more flavor, and/or more final-state partons. I.e., they can exploit much higher cross sections and/or more complicated event topologies. With the search that we are proposing, we have finally “hit bottom” on colored particle pair-production and decay. *Any* model with a color-triplet “diquark” scalar can be searched for, whether it is connected to naturalness or not, and independently of the flavor structure of its decay (see e.g. [70–73]).¹¹

¹¹The other choice of diquark color representation, namely the color-sextet, has not been explicitly searched for, but its cross section is actually larger than the color-octet. We therefore expect that the limits on color-sextet scalars with prompt two-body decays are more constraining than those for color-octets [12, 13].

In terms of concrete performance, we have found that jet substructure at the LHC can conclusively push beyond the 100 GeV threshold set by previous limits, leaving no gaps. In fact, it should be possible to achieve discovery-level sensitivity at 100 GeV using 2012 data, demonstrating the LHC’s far superior production cross sections and luminosity compared to earlier experiments. Exclusion-level sensitivity extending up to 300 GeV or higher seems achievable. These results can be obtained using a number of promising data-driven search techniques, and should be realistic if at least one of these techniques exhibits manageable systematic errors. Returning to the usual MFV-inspired assumption that most decays contain b -quarks, the option of an extremely powerful b -tagged search opens up. This may provide roughly 10σ discovery sensitivity at 100 GeV, and exclusions extending to nearly 400 GeV, without accounting for possible re-optimization of the analysis cuts or further gains from applying a double- b -tag. The upcoming Run II of the LHC should push exclusion to masses about twice as high, even while further tightening sensitivity at lower masses.

These findings have yet again illustrated the utility of viewing complicated hadronic activity through the lens of jet substructure. This basic change in philosophy from canonical jet-based reconstruction (almost) frees us of the notion of a “minimum distance” between reconstructable hard partons, namely the jet radius, which was in fact a major bottleneck in the traditional 4-jet versions of this search. Clearly with a more flexible viewpoint, the extremely high energies available at the LHC can be made to work for us, not against us, when attempting to search for $O(100\text{ GeV})$ or even $O(10\text{'s of GeV})$ objects decaying into jets. Indeed, we have found here that we are in a good position to test whether supersymmetry has exploited the limitations of our conventional analyses, and has been hiding in plain sight this whole time.

Acknowledgments

We thank Jared Evans, Roberto Franceschini, David Miller, and Keith Ulmer for useful discussions and comments. We thank Stefan Prestel for much help in using CKKW-L matching and for fixing a bug in `Pythia8 v1.7.6`. Matched QCD simulations were performed on the Harvard Odyssey Cluster. YB was supported by start-up funds from the Univ. of Wisconsin, Madison. AK was supported by NSF grant No. PHY-0855591. BT was supported by DoE grant Nos. DE-FG-02-91ER40676 and DE-FG-02-95ER40896, by NSF grant No. PHY-0969510 (LHC Theory Initiative), and by the PITT PACC. YB would also like to thank the Kavli Institute for Theoretical Physics, U.C. Santa Barbara, which is supported under Grant No. NSF PHY11-25915. AK is also grateful to the Galileo Galilei Institute in Florence where this work was partially done. We also thank the Aspen Center for Physics, under NSF Grant No. PHY-1066293, where this work was both initiated and completed.

A Simulation details

Our nominal stop pair simulations begin in `MadGraph5 v1.4.7` [74] with a UFO simplified model [75] containing stops as the only BSM particle. The simulations are matched up to one jet emission in production using the default `PYTHIA6` [76]. The matching is of the

k_T -MLM type, with a matching scale of 30 GeV. There is a technical difficulty in such a simulation, in that `PYTHIA6` does not recognize the color epsilon-tensor structure appearing in the stop's RPV two-body decay. To bypass this difficulty, we first hadronize the stops, along with the rest of the event, using publicly-available R-hadronization code [77].¹² We then simulate the stop-hadron decays by treating them as if they were Z^* s of the same mass, decaying to $b\bar{d}$ and showering/hadronizing independently of the rest of the event.¹³ This should build up a decay showering pattern essentially identical to what would have been obtained in `Pythia8` [78].

It is also possible to simulate stop pair events with RPV decays directly in `Pythia8`, for which we use `v1.7.6` [79]. Our initial simulations indicated significant disagreements in the modeling of the production-stage radiation of high- $p_T(\tilde{t})$ events relative to both matched `MadGraph5` and `Prospino 1.0` [80], and spurious m_{avg} tails in our final reconstructions. Subsequently, we found that this behavior is driven by the default “power shower” treatment of events without light partons in the final state, and can be made to yield a good fit to the matrix element predictions by using shower damping (`SpaceShower:pTdampMatch=1`). (By contrast, we have found that the “wimpy shower,” as well as the virtuality-ordered and p_T -ordered showers of `PYTHIA6`, all tend to underestimate the amount of production radiation.) Using self-contained `Pythia8` simulations with shower damping, and running through our complete analysis chain described below, we obtain very good agreement with the more complicated matched simulations. Running matching within `Pythia8` is also possible, but we have not explored this option for our signal generation.

Our nominal QCD continuum simulations are based on `MadGraph5` matched up to four partons within `Pythia8`, using the more theoretically-rigorous CKKW-L prescription [61].¹⁴ For our merging measure, we use the Durham- k_T distance used internally by `MadGraph5`,¹⁵ and exploit the program's ability to produce multi-parton simulations with kinematic cuts defined in that space with threshold `xqcut`.¹⁶ The measure only applies to partons that can realistically be viewed as merging within a QCD diagram according to quark flavor. For example, two antiquarks would never be compared, nor would a u -quark and a c -antiquark. We also forbid mergings with the beam if they would violate flavor, though this occurrence is very rare. Our merging scale is set to 50 GeV for 8 TeV simulations, and 100 GeV for 14 TeV simulations. (For the former, the m_{avg} range over which we run our stop searches is highly dominated by the hard 4-parton events.)

¹²When running with multi-parton interactions (i.e., underlying event), we found that this code runs much more seamlessly with the virtuality-ordered shower, and have used this for our simulations.

¹³It is worth pointing out that, over a large range of acceptable parameters for the RPV couplings, the stops should indeed hadronize before decaying. In any case, we expect that the order-of-operations of decay versus hadronization will only have a very minor impact on the final radiation pattern as seen by our procedures.

¹⁴`v1.7.6` of `Pythia8` is affected by a bug that causes crashes for this matching. We thank Stefan Prestel for providing us with a private fix.

¹⁵Beam-distances are computed as $k_{T,iB}^2 = p_{T,i}^2 + m_i^2$. Inter-parton distances are computed as $k_{T,ij}^2 = \max(m_i^2, m_j^2) + 2 \min(p_{T,i}^2, p_{T,j}^2) \times (\cosh \Delta\eta_{ij} - \cos \Delta\phi_{ij})$.

¹⁶To avoid double-counting of α_s reweightings, we have commented-out the relevant code in `MadGraph5`.

For comparison, we have also generated a set of unmatched QCD simulations wholly within `Pythia8`, based on showered $2 \rightarrow 2$ production with the default (wimpy) shower. These display a qualitatively similar m_{avg} spectrum to the matched sample, though with a less steep falloff, and $O(1)$ higher rate at $m_{\text{avg}} \gtrsim 200$ GeV. These differences are much less pronounced when comparing samples processed through ordinary jet reconstruction.

For our other backgrounds ($t\bar{t}$, W/Z +jets, t/s -channel single-top, tW , diboson), we mainly relied on `Pythia8`, though our $t\bar{t}$ simulation starts as a $2 \rightarrow 6$ decay chain processes in `MadGraph5` in order to capture spin correlation effects. We have also damped the shower for $t\bar{t}$, which is treated as a power shower by default. Similar to stop pair production, comparison to the kinematics of undecayed matched samples shows good agreement.

All of our simulations are leading-order, and can be normalized to higher-precision calculations using K-factors. For our matched stop samples, we find that a flat K-factor of 1.5 corrects us to the NLO+NLL predictions [81] for almost any mass. We have also verified, by comparing with $p_T(\tilde{t})$ spectra predicted by `Prospino 1.0`, that the matched spectrum is in excellent agreement with the NLO spectrum. Therefore we do not anticipate a substantially different K-factor for boosted stops. For $t\bar{t}$, we use a K-factor of 1.8. For all of our other simulations, including matched QCD, we coarsely assume K-factors of 1.5. Our own comparisons of matched and unmatched QCD simulations to the data of [11] suggest that this choice of K-factor may be conservatively large.

Downstream of the hadron-level simulation, we apply a simple detector model in the form of a 0.1×0.1 calorimeter grid in η - ϕ space. We form our final H_T trigger by first clustering the cells into $R = 0.5$ anti- k_T jets with `FastJet3` [59], and demanding that the sum of p_T s of jets above 50 GeV exceeds 900 GeV (1600 GeV for 14 TeV simulations). At this stage we neglect both jet energy measurement fluctuations and pileup effects. The former should be only a few percent for such a large sum-over-energies, and the latter would be systematically accounted for in a realistic jet measurement.

While we do not concern ourselves with the effects of pileup on ordinary jet energy measurements, we do wish to make sure that our subsequent jet substructure methods are not adversely affected. Therefore, on top of each hard event surviving the H_T trigger, we superimpose an average of 20 min-bias events from `Pythia8`, including both charged and neutral activity. We then actively remove this pileup using a slightly modified form of trimming [63]. We cluster all calorimeter cells in the event into $R = 0.2$ anti- k_T jets, and discard the contents of these jets if their total p_T falls below an absolute cutoff of 5 GeV. Unlike canonical trimming, which works jet-by-jet and operates with a *relative* p_T measure, we have targeted this method to subtract the contaminating energy density of pileup, at least in regions that do not overlap with hard activity. We have found that this procedure successfully preserves our 100 GeV stop lineshape, which is otherwise shifted by 5 GeV and broadened due to the pileup. The stop reconstruction rate is unchanged. The final impact on the background is generally modest, though there is an $O(1)$ reduction of the high- m_{avg} tail, e.g. in the vicinity of 300 GeV.

Following this pre-trimming stage, we apply our main clustering into $R = 1.5$ C/A fat-jets. Only the leading two fat-jets in p_T are considered, and these must both fall in the region $|\eta| < 2.5$. The fat-jets are declustered into subjets as described in section 2. To




Figure 7. The ΔR distributions of subjects within reconstructed fat-jets passing all analysis cuts, for stops of mass 100 GeV (blue), 200 GeV (purple), 300 GeV (red), and 400 GeV (green). (Small spikes at $\Delta R = 0.1$ correspond to events where both stop-jets have been declustered down to our calorimeter model granularity, and would have $m_{\text{avg}} \sim 10$ GeV.)

obtain more realistic stop mass peaks, we smear the energies of the subjects as¹⁷

$$\frac{\Delta E}{E} = \frac{5 \text{ GeV}}{E} \oplus \frac{0.5 \text{ GeV}^{1/2}}{\sqrt{E}} \oplus 0.05 . \quad (\text{A.1})$$

To cover scenarios where stops have large branching fractions to b -quarks, we also run a b -tagged analysis. We assume that subjects can be tagged similar to ordinary jets, and that the tag/mistag efficiencies are fairly flat in our analysis range. (Our subject p_T 's typically vary between 100 GeV and 400 GeV. See, e.g., [82].) To perform flavor tagging, we keep track of bottom-hadrons and prompt charm-hadrons from the event record, and match them to the closest subject within $\Delta R < 0.2$. Each subject's “true” flavor is then determined by the heaviest associated hadron. We apply flat b -tagging efficiencies of 60%, 10%, and 2% for bottom-flavored, charm-flavored, and unflavored subjects, respectively.

B Supplementary results

This appendix contains three supplementary sets of results: the ΔR distributions of subjects for signal events, a comparison of our nominal $R = 1.5$ jet radius to $R = 0.8$, and comparisons with the more standard BDRS declustering procedure.

Figure 7 shows the ΔR distributions of subjects within stop-jets, for events passing our complete set of analysis cuts. This plot makes it clear that for $m_{\tilde{t}} = 100$ GeV, a large fraction of stop decays would comfortably sit inside of a normal-sized LHC jet of $R = 0.4$

¹⁷More ideally, we would smear the energies of individual calorimeter cells, which would also allow us to simulate how imperfections in the energy measurements affect the declustering. However, as we do not expect to be able to model this well, we still treat the energy measurements as perfect at that stage. Still, our final analysis cut on the p_T ratio of the subjects operates on smeared objects.




Figure 8. Spectra in m_{avg} for matched QCD passing all cuts, reconstructed with the nominal $R = 1.5$ fat-jets (black) or with $R = 0.8$ (blue).

or $R = 0.5$. It is also notable that, even though we choose a much larger fat-jet radius, very few stop decays are reconstructed with unphysically-large ΔR . In other words, our substructure procedures and analysis cuts adaptively find the “correct” ΔR scale for the signal. For larger stop masses, the separation becomes large enough that an ordinary jet radius could resolve the decays. But in our treatment this regime is continuously connected to the scenarios with $\Delta R < 0.4$, with no artificial threshold. Finally, we can see that with our absolute and relative energy cuts, $m_{\tilde{t}} = 300$ GeV is about the largest mass that displays complete containment within $R = 1.5$ fat-jets. Still, a large fraction of $m_{\tilde{t}} = 400$ GeV decays remain contained, a signal which is important for the b -tagged version of the analysis.

In figure 8, we compare the QCD continuum’s m_{avg} spectrum with our nominal $R = 1.5$ to an identical analysis with $R = 0.8$. It can be seen that, in the vicinity of $m_{\text{avg}} = 100$ GeV, the background increases both in absolute rate and in steepness. Essentially, the entire spectrum has been “squashed” by a factor of 2, since the overall mass scale is set by $R \times H_T$. Performing the same analysis with the $m_{\tilde{t}} = 100$ GeV signal, the lineshape is practically unaltered, but the overall acceptance increases by 30%. This is because, with a narrower fat-jet, there are fewer cases where the declustering picks up a spurious ISR jet. Still, the gain in S/\sqrt{B} is marginal, and comes at a cost of slightly reduced S/B in addition to a more difficult background shape. Higher stop masses display significantly reduced efficiencies due to incomplete containment.

Our nominal declustering procedure judges splittings based on the p_T ’s of subjet candidates relative to the original fat-jet, and on their individual m/p_T ratios. This procedure is a direct descendant of the BDRS procedure of [43], which uses a somewhat different set of declustering criteria, and also applies an additional filtering step by reclustering the subjet constituents. Within the kinematic regime of our analysis, the declustering stage of BDRS acts almost identically to our procedure without the m/p_T requirement.¹⁸ With

¹⁸The BDRS mass-drop criterion is mostly redundant here, and the declustering is driven mainly by the momentum-asymmetry criterion. See [83] for a detailed related study.

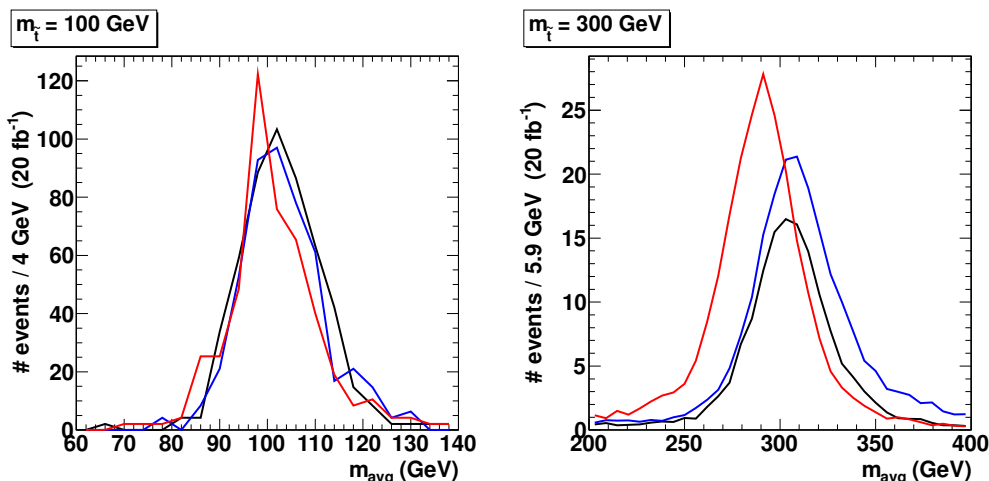


Figure 9. Stop signal peak reconstructions for $m_{\tilde{t}} = 100$ GeV (left) and $m_{\tilde{t}} = 300$ GeV (right), using our nominal relative- p_T declustering (black), full BDRS with filtering (red), and BDRS without filtering (blue). (No pileup or trimming have been applied.)

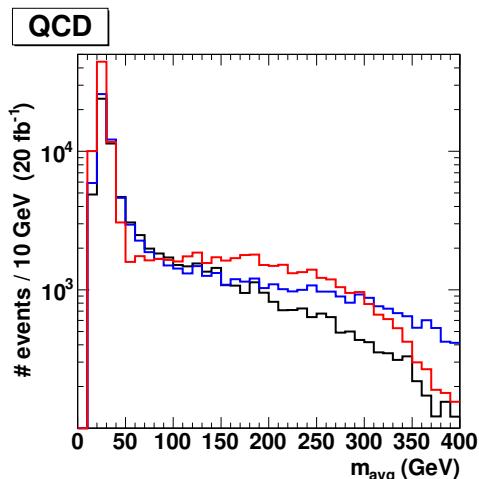


Figure 10. Matched QCD reconstructions, using our nominal relative- p_T declustering (black), full BDRS with filtering (red), and BDRS without filtering (blue). (No pileup or trimming have been applied.)

filtering, the two subjets are further refined into three, using the C/A algorithm with $R = \min[\Delta R(\text{subjets})/2, 0.3]$. We reform these back into two subjets, by clustering together the two that are closest in ΔR . This allows us to apply our final cut on the ratio of subjet p_T 's, which assumes 2-body substructure. We also consider a form of BDRS without the filtering step. Figures 9 and 10 show a comparison of our nominal procedure with both filtered and unfiltered BDRS. (This comparison is made without pileup or trimming.) It is clear that the 100 GeV stop signal is fairly insensitive to the detailed procedure, but that the 300 GeV stop and QCD spectra can be highly reshaped, and that the overall rates in the vicinity of $m_{\text{avg}} = 300$ GeV are increased by $O(1)$. The filtered QCD spectrum is also

flatter in the region between 100 GeV and 200 GeV, which could help a signal bump stand out more clearly there. These last two points can be considered advantages. However, filtering also accidentally introduces new mass scales from its minimum reclustering radius of 0.3 for well-separated subjects, resulting in a bimodal QCD distribution with a local minimum near 50 GeV and a broad local maximum near 175 GeV. And while the 300 GeV stop signal is enhanced, it sits on top of a background that is starting to sharply change shape. Therefore, these advantages must be treated with caution. Removing filtering, we still see some enhancement of the 300 GeV signal and its background, but the background shape becomes less biased. The differences between unfiltered BDRS and our nominal declustering are dominated by the introduction of high- m/p_T subjects for the former, and these are more likely to be contaminated by additional radiation. We have also found that the detailed shape of the spectrum at high mass develops more sensitivity to changing analysis cuts to define control regions. However, these differences relative our nominal procedure might be reduced with the use of trimming.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] ATLAS collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Phys. Lett. B* **716** (2012) 1 [[arXiv:1207.7214](https://arxiv.org/abs/1207.7214)] [[INSPIRE](#)].
- [2] CMS collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30 [[arXiv:1207.7235](https://arxiv.org/abs/1207.7235)] [[INSPIRE](#)].
- [3] S. Dimopoulos and G. Giudice, *Naturalness constraints in supersymmetric theories with nonuniversal soft terms*, *Phys. Lett. B* **357** (1995) 573 [[hep-ph/9507282](https://arxiv.org/abs/hep-ph/9507282)] [[INSPIRE](#)].
- [4] A.G. Cohen, D. Kaplan and A. Nelson, *The more minimal supersymmetric Standard Model*, *Phys. Lett. B* **388** (1996) 588 [[hep-ph/9607394](https://arxiv.org/abs/hep-ph/9607394)] [[INSPIRE](#)].
- [5] Y. Kats, P. Meade, M. Reece and D. Shih, *The status of GMSB after 1 fb⁻¹ at the LHC*, *JHEP* **02** (2012) 115 [[arXiv:1110.6444](https://arxiv.org/abs/1110.6444)] [[INSPIRE](#)].
- [6] C. Brust, A. Katz, S. Lawrence and R. Sundrum, *SUSY, the third generation and the LHC*, *JHEP* **03** (2012) 103 [[arXiv:1110.6670](https://arxiv.org/abs/1110.6670)] [[INSPIRE](#)].
- [7] M. Papucci, J.T. Ruderman and A. Weiler, *Natural SUSY endures*, *JHEP* **09** (2012) 035 [[arXiv:1110.6926](https://arxiv.org/abs/1110.6926)] [[INSPIRE](#)].
- [8] C. Brust, A. Katz and R. Sundrum, *SUSY stops at a bump*, *JHEP* **08** (2012) 059 [[arXiv:1206.2353](https://arxiv.org/abs/1206.2353)] [[INSPIRE](#)].
- [9] R. Barbier et al., *R-parity violating supersymmetry*, *Phys. Rept.* **420** (2005) 1 [[hep-ph/0406039](https://arxiv.org/abs/hep-ph/0406039)] [[INSPIRE](#)].
- [10] CMS collaboration, *Search for new physics in the paired dijet mass spectrum*, [CMS-PAS-EXO-11-016](https://arxiv.org/abs/1109.4548), CERN, Geneva Switzerland (2011).

- [11] CMS collaboration, *Search for pair-produced dijet resonances in four-jet final states in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev. Lett.* **110** (2013) 141802 [[arXiv:1302.0531](#)] [[INSPIRE](#)].
- [12] ATLAS collaboration, *Search for massive colored scalars in four-jet final states in $\sqrt{s} = 7$ TeV proton-proton collisions with the ATLAS detector*, *Eur. Phys. J. C* **71** (2011) 1828 [[arXiv:1110.2693](#)] [[INSPIRE](#)].
- [13] ATLAS collaboration, *Search for pair-produced massive coloured scalars in four-jet final states with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **73** (2013) 2263 [[arXiv:1210.4826](#)] [[INSPIRE](#)].
- [14] J.A. Evans and Y. Kats, *LHC coverage of RPV MSSM with light stops*, *JHEP* **04** (2013) 028 [[arXiv:1209.0764](#)] [[INSPIRE](#)].
- [15] C. Csáki, Y. Grossman and B. Heidenreich, *MFV SUSY: a natural theory for R-parity violation*, *Phys. Rev. D* **85** (2012) 095009 [[arXiv:1111.1239](#)] [[INSPIRE](#)].
- [16] P.W. Graham, D.E. Kaplan, S. Rajendran and P. Saraswat, *Displaced supersymmetry*, *JHEP* **07** (2012) 149 [[arXiv:1204.6038](#)] [[INSPIRE](#)].
- [17] G. Krnjaic and D. Stolarski, *Gauging the way to MFV*, *JHEP* **04** (2013) 064 [[arXiv:1212.4860](#)] [[INSPIRE](#)].
- [18] B. Bhattacharjee, J.L. Evans, M. Ibe, S. Matsumoto and T.T. Yanagida, *Natural SUSY's last hope: R-parity violation via UDD operators*, *Phys. Rev. D* **87** (2013) 115002 [[arXiv:1301.2336](#)] [[INSPIRE](#)].
- [19] R. Franceschini and R. Mohapatra, *New patterns of natural R-parity violation with supersymmetric gauged flavor*, *JHEP* **04** (2013) 098 [[arXiv:1301.3637](#)] [[INSPIRE](#)].
- [20] C. Csáki and B. Heidenreich, *A complete model for R-parity violation*, *Phys. Rev. D* **88** (2013) 055023 [[arXiv:1302.0004](#)] [[INSPIRE](#)].
- [21] A. Florez, D. Restrepo, M. Velasquez and O. Zapata, *Baryonic violation of R-parity from anomalous $U(1)_H$* , *Phys. Rev. D* **87** (2013) 095010 [[arXiv:1303.0278](#)] [[INSPIRE](#)].
- [22] G. Krnjaic and Y. Tsai, *Soft RPV through the baryon portal*, [arXiv:1304.7004](#) [[INSPIRE](#)].
- [23] A. Monteux, *Natural, R-parity violating supersymmetry and horizontal flavor symmetries*, *Phys. Rev. D* **88** (2013) 045029 [[arXiv:1305.2921](#)] [[INSPIRE](#)].
- [24] L. Di Luzio, M. Nardecchia and A. Romanino, *A framework for baryonic R-parity violation in grand unified theories*, *Phys. Rev. D* **88** (2013) 115008 [[arXiv:1305.7034](#)] [[INSPIRE](#)].
- [25] C. Csáki, E. Kuflik and T. Volansky, *Dynamical R-parity violation*, [arXiv:1309.5957](#) [[INSPIRE](#)].
- [26] CMS collaboration, *Search for top squarks in R-parity-violating supersymmetry using three or more leptons and b-tagged jets*, *Phys. Rev. Lett.* **111** (2013) 221801 [[arXiv:1306.6643](#)] [[INSPIRE](#)].
- [27] ATLAS collaboration, *Search for supersymmetry in events with four or more leptons in 21 fb^{-1} of pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, ATLAS-CONF-2013-036, CERN, Geneva Switzerland (2013).
- [28] ATLAS collaboration, *Search for massive particles in multijet signatures with the ATLAS detector in $\sqrt{s} = 8$ TeV pp collisions at the LHC*, ATLAS-CONF-2013-091, CERN, Geneva Switzerland (2013).

- [29] ATLAS collaboration, *Search for pair production of massive particles decaying into three quarks with the ATLAS detector in $\sqrt{s} = 7$ TeV pp collisions at the LHC*, *JHEP* **12** (2012) 086 [[arXiv:1210.4813](#)] [[INSPIRE](#)].
- [30] CMS collaboration, *Search for light- and heavy-flavor three-jet resonances in multijet final states at 8 TeV*, *CMS-PAS-EXO-12-049*, CERN, Geneva Switzerland (2012).
- [31] CMS collaboration, *Search for three-jet resonances in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Lett. B* **718** (2012) 329 [[arXiv:1208.2931](#)] [[INSPIRE](#)].
- [32] ALEPH collaboration, A. Heister et al., *Search for supersymmetric particles with R -parity violating decays in e^+e^- collisions at \sqrt{s} up to 209 GeV*, *Eur. Phys. J. C* **31** (2003) 1 [[hep-ex/0210014](#)] [[INSPIRE](#)].
- [33] CDF collaboration, T. Aaltonen et al., *Search for pair-production of strongly-interacting particles decaying to pairs of jets in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Rev. Lett.* **111** (2013) 031802 [[arXiv:1303.2699](#)] [[INSPIRE](#)].
- [34] D. Duggan et al., *Sensitivity of an upgraded LHC to R -parity violating signatures of the MSSM*, [arXiv:1308.3903](#) [[INSPIRE](#)].
- [35] Z. Han, A. Katz, M. Son and B. Tweedie, *Boosting searches for natural SUSY with RPV via gluino cascades*, *Phys. Rev. D* **87** (2013) 075003 [[arXiv:1211.4025](#)] [[INSPIRE](#)].
- [36] M. Lisanti, P. Schuster, M. Strassler and N. Toro, *Study of LHC searches for a lepton and many jets*, *JHEP* **11** (2012) 081 [[arXiv:1107.5055](#)] [[INSPIRE](#)].
- [37] B. Allanach and B. Gripaios, *Hide and seek with natural supersymmetry at the LHC*, *JHEP* **05** (2012) 062 [[arXiv:1202.6616](#)] [[INSPIRE](#)].
- [38] J. Berger, M. Perelstein, M. Saelim and P. Tanedo, *The same-sign dilepton signature of RPV/MFV SUSY*, *JHEP* **04** (2013) 077 [[arXiv:1302.2146](#)] [[INSPIRE](#)].
- [39] ATLAS collaboration, *Search for strongly produced superpartners in final states with two same sign leptons with the ATLAS detector using 21 fb^{-1} of proton-proton collisions at $\sqrt{s} = 8$ TeV*, *ATLAS-CONF-2013-007*, CERN, Geneva Switzerland (2013).
- [40] R. Franceschini and R. Torre, *RPV stops bump off the background*, *Eur. Phys. J. C* **73** (2013) 2422 [[arXiv:1212.3622](#)] [[INSPIRE](#)].
- [41] M.H. Seymour, *Searches for new particles using cone and cluster jet algorithms: a comparative study*, *Z. Phys. C* **62** (1994) 127 [[INSPIRE](#)].
- [42] J. Butterworth, B. Cox and J.R. Forshaw, *WW scattering at the CERN LHC*, *Phys. Rev. D* **65** (2002) 096014 [[hep-ph/0201098](#)] [[INSPIRE](#)].
- [43] J.M. Butterworth, A.R. Davison, M. Rubin and G.P. Salam, *Jet substructure as a new Higgs search channel at the LHC*, *Phys. Rev. Lett.* **100** (2008) 242001 [[arXiv:0802.2470](#)] [[INSPIRE](#)].
- [44] NEW PHYSICS WORKING GROUP collaboration, G. Brooijmans et al., *New physics at the LHC. A Les Houches report: physics at TeV colliders 2009 — new physics working group*, [arXiv:1005.1229](#) [[INSPIRE](#)].
- [45] D. Curtin, R. Essig and B. Shuve, *Boosted multijet resonances and new color-flow variables*, *Phys. Rev. D* **88** (2013) 034019 [[arXiv:1210.5523](#)] [[INSPIRE](#)].
- [46] Y. Bai and J. Shelton, *Composite octet searches with jet substructure*, *JHEP* **07** (2012) 067 [[arXiv:1107.3563](#)] [[INSPIRE](#)].

- [47] J.M. Butterworth, J.R. Ellis, A.R. Raklev and G.P. Salam, *Discovering baryon-number violating neutralino decays at the LHC*, *Phys. Rev. Lett.* **103** (2009) 241803 [[arXiv:0906.0728](#)] [[INSPIRE](#)].
- [48] A. Katz, M. Son and B. Tweedie, *Jet substructure and the search for neutral spin-one resonances in electroweak boson channels*, *JHEP* **03** (2011) 011 [[arXiv:1010.5253](#)] [[INSPIRE](#)].
- [49] M. Gouzevitch et al., *Scale-invariant resonance tagging in multijet events and new physics in Higgs pair production*, *JHEP* **07** (2013) 148 [[arXiv:1303.6636](#)] [[INSPIRE](#)].
- [50] D.E. Kaplan, K. Rehermann, M.D. Schwartz and B. Tweedie, *Top tagging: a method for identifying boosted hadronically decaying top quarks*, *Phys. Rev. Lett.* **101** (2008) 142001 [[arXiv:0806.0848](#)] [[INSPIRE](#)].
- [51] M. Son, C. Spethmann and B. Tweedie, *Diboson-jets and the search for resonant Zh production*, *JHEP* **08** (2012) 160 [[arXiv:1204.0525](#)] [[INSPIRE](#)].
- [52] S.D. Ellis, C.K. Vermilion and J.R. Walsh, *Techniques for improved heavy particle searches with jet substructure*, *Phys. Rev. D* **80** (2009) 051501 [[arXiv:0903.5081](#)] [[INSPIRE](#)].
- [53] J. Thaler and K. Van Tilburg, *Identifying boosted objects with N-subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [54] L.G. Almeida, S.J. Lee, G. Perez, G. Sterman and I. Sung, *Template overlap method for massive jets*, *Phys. Rev. D* **82** (2010) 054034 [[arXiv:1006.2035](#)] [[INSPIRE](#)].
- [55] D.E. Soper and M. Spannowsky, *Finding physics signals with shower deconstruction*, *Phys. Rev. D* **84** (2011) 074002 [[arXiv:1102.3480](#)] [[INSPIRE](#)].
- [56] S.D. Ellis, A. Hornig, T.S. Roy, D. Krohn and M.D. Schwartz, *Qjets: a non-deterministic approach to tree-based jet substructure*, *Phys. Rev. Lett.* **108** (2012) 182003 [[arXiv:1201.1914](#)] [[INSPIRE](#)].
- [57] Y.L. Dokshitzer, G. Leder, S. Moretti and B. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)] [[INSPIRE](#)].
- [58] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, [hep-ph/9907280](#) [[INSPIRE](#)].
- [59] M. Cacciari and G.P. Salam, *Dispelling the N^3 myth for the k_t jet-finder*, *Phys. Lett. B* **641** (2006) 57 [[hep-ph/0512210](#)] [[INSPIRE](#)].
- [60] CMS collaboration, *Search for supersymmetry in hadronic final states using M_{T2} in pp collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **10** (2012) 018 [[arXiv:1207.1798](#)] [[INSPIRE](#)].
- [61] L. Lönnblad, *Correcting the color dipole cascade model with fixed order matrix elements*, *JHEP* **05** (2002) 046 [[hep-ph/0112284](#)] [[INSPIRE](#)].
- [62] L. Lönnblad and S. Prestel, *Matching tree-level matrix elements with interleaved showers*, *JHEP* **03** (2012) 019 [[arXiv:1109.4829](#)] [[INSPIRE](#)].
- [63] D. Krohn, J. Thaler and L.-T. Wang, *Jet trimming*, *JHEP* **02** (2010) 084 [[arXiv:0912.1342](#)] [[INSPIRE](#)].
- [64] T. Plehn, M. Spannowsky, M. Takeuchi and D. Zerwas, *Stop reconstruction with tagged tops*, *JHEP* **10** (2010) 078 [[arXiv:1006.2833](#)] [[INSPIRE](#)].

- [65] ATLAS collaboration, *Measurements of the properties of the Higgs-like boson in the two photon decay channel with the ATLAS detector using 25 fb⁻¹ of proton-proton collision data*, [ATLAS-CONF-2013-012](#), CERN, Geneva Switzerland (2013).
- [66] CMS collaboration, *Properties of the observed Higgs-like resonance using the diphoton channel*, [CMS-PAS-HIG-13-016](#), CERN, Geneva Switzerland (2013).
- [67] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [[arXiv:1307.0007](#)] [[INSPIRE](#)].
- [68] M. Dasgupta, A. Fregoso, S. Marzani and A. Powling, *Jet substructure with analytical methods*, *Eur. Phys. J. C* **73** (2013) 2623 [[arXiv:1307.0013](#)] [[INSPIRE](#)].
- [69] S. El Hedri, A. Hook, M. Jankowiak and J.G. Wacker, *Learning how to count: a high multiplicity search for the LHC*, *JHEP* **08** (2013) 136 [[arXiv:1302.1870](#)] [[INSPIRE](#)].
- [70] E. Ma, M. Raidal and U. Sarkar, *Probing the exotic particle content beyond the Standard Model*, *Eur. Phys. J. C* **8** (1999) 301 [[hep-ph/9808484](#)] [[INSPIRE](#)].
- [71] E. Del Nobile, R. Franceschini, D. Pappadopulo and A. Strumia, *Minimal matter at the Large Hadron Collider*, *Nucl. Phys. B* **826** (2010) 217 [[arXiv:0908.1567](#)] [[INSPIRE](#)].
- [72] G.F. Giudice, B. Gripaios and R. Sundrum, *Flavourful production at hadron colliders*, *JHEP* **08** (2011) 055 [[arXiv:1105.3161](#)] [[INSPIRE](#)].
- [73] L. Vecchi, *Color & weak triplet scalars, the dimuon asymmetry in B_s decay, the top forward-backward asymmetry and the CDF dijet excess*, *JHEP* **10** (2011) 003 [[arXiv:1107.2933](#)] [[INSPIRE](#)].
- [74] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer and T. Stelzer, *MadGraph 5: going beyond*, *JHEP* **06** (2011) 128 [[arXiv:1106.0522](#)] [[INSPIRE](#)].
- [75] C. Degrande et al., *UFO — the Universal FeynRules output*, *Comput. Phys. Commun.* **183** (2012) 1201 [[arXiv:1108.2040](#)] [[INSPIRE](#)].
- [76] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].
- [77] *PYTHIA example main programs webpage*, <http://pythia6.hepforge.org/examples/>.
- [78] N. Desai and P.Z. Skands, *Supersymmetry and generic BSM models in PYTHIA 8*, *Eur. Phys. J. C* **72** (2012) 2238 [[arXiv:1109.5852](#)] [[INSPIRE](#)].
- [79] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [80] W. Beenakker, R. Hopker and M. Spira, *PROSPINO: a program for the production of supersymmetric particles in next-to-leading order QCD*, [hep-ph/9611232](#) [[INSPIRE](#)].
- [81] W. Beenakker et al., *Supersymmetric top and bottom squark production at hadron colliders*, *JHEP* **08** (2010) 098 [[arXiv:1006.4771](#)] [[INSPIRE](#)].
- [82] CMS collaboration, *Performance of b tagging at $\sqrt{s} = 8$ TeV in multijet, $t\bar{t}$ and boosted topology events*, [CMS-PAS-BTV-13-001](#), CERN, Geneva Switzerland (2013).
- [83] Y. Cui, Z. Han and M.D. Schwartz, *W-jet tagging: optimizing the identification of boosted hadronically-decaying W bosons*, *Phys. Rev. D* **83** (2011) 074023 [[arXiv:1012.2077](#)] [[INSPIRE](#)].