

Sequential displaced vertices: Novel collider signature for long-lived particles

Keith R. Dienes^{1,2,*}, Doojin Kim^{3,†}, Tara T. Leininger^{4,‡} and Brooks Thomas^{4,§}

¹*Department of Physics, University of Arizona, Tucson, Arizona 85721, USA*

²*Department of Physics, University of Maryland, College Park, Maryland 20742, USA*

³*Mitchell Institute for Fundamental Physics and Astronomy, Department of Physics and Astronomy,
Texas A&M University, College Station, Texas 77843 USA*

⁴*Department of Physics, Lafayette College, Easton, Pennsylvania 18042, USA*



(Received 13 August 2021; accepted 11 October 2022; published 9 November 2022)

In this paper, we point out a novel signature of physics beyond the Standard Model which could potentially be observed both at the Large Hadron Collider (LHC) and at future colliders. This signature, which emerges naturally within many proposed extensions of the Standard Model, results from the multiple displaced vertices associated with the successive decays of unstable, long-lived particles along the same decay chain. We call such a sequence of displaced vertices a “tumbler.” We examine the prospects for observing tumblers at the LHC and assess the extent to which tumbler signatures can be distinguished from other signatures of new physics which also involve multiple displaced vertices within the same collider event. As part of this analysis, we also develop a procedure for reconstructing the masses and lifetimes of the particles involved in the corresponding decay chains. We find that the prospects for discovering and distinguishing tumblers can be greatly enhanced by exploiting precision timing information such as would be provided by the CMS timing layer at the high-luminosity LHC. Our analysis therefore provides strong additional motivation for continued efforts to improve the timing capabilities of collider detectors at the LHC and beyond.

DOI: [10.1103/PhysRevD.106.095012](https://doi.org/10.1103/PhysRevD.106.095012)

I. INTRODUCTION

Ever since the seminal work of Glashow, Weinberg, and Salam in the 1970s that gave birth to modern particle physics, the Standard Model (SM) has reigned supreme. Although the discovery of neutrino oscillations and the preponderance of observational evidence for dark matter and dark energy have indicated the need to extend the SM into new domains, the core of the SM has remained intact and continues to accurately describe all existing collider data despite decades of intense experimental research. Indeed, unambiguous evidence for possible SM extensions such as weak-scale supersymmetry or large extra dimensions has not yet been found.

There are, in principle, two possible reasons for this state of affairs. On the one hand, the energy scale

associated with the new physics may be sufficiently high that this physics lies beyond the reach of current experiments. However, on the other hand, it is possible that the new physics resides at energy scales which are potentially accessible at current or imminent collider experiments, but that this physics is manifested through collider signatures that have not yet received much attention within the community (for reviews, see, e.g., Refs. [1–3]).

In this paper we point out a novel collider signature which arises in a variety of scenarios for new physics. This signature rests on the possible existence of long-lived particles (LLPs). As discussed in Ref. [4], LLPs can arise in many proposed extensions of the Standard Model. These include models which attempt to address the gauge hierarchy problem, models which provide new approaches to dark-matter physics, models which describe different scenarios for baryogenesis and leptogenesis, and even nonminimal models of neutrino physics. Because of their relative long lifetimes, LLPs, once produced, can propagate across macroscopic distances before they decay. For LLPs with proper decay lengths $c\tau$ ranging from millimeters to hundreds of meters, these decays can give rise to a number of distinctive signatures at colliders, including emerging jets [5], disappearing tracks, and macroscopically displaced vertices (DVs). While searches for DVs are part of the standard experimental program at colliders, the signature

*dienes@arizona.edu

†doojin.kim@tamu.edu

‡leiningt@lafayette.edu

§thomasbd@lafayette.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

on which we shall focus our attention involves the presence of multiple displaced vertices which result from the successive decays of multiple unstable LLPs within the same decay chain. In such cases, the event unfolds by “tumbling” down the steps of the decay chain, terminating only once a collider-stable particle is reached.

Given this decay topology, we shall refer to such a sequence of DVs as a “tumbler.” In this work, we shall consider the special case of tumblers in which each such LLP decay yields a single, lighter LLP as well as one or more SM particles which can be detected directly by a collider detector. The signatures of such tumblers are quite striking as they have very low SM backgrounds. We shall examine the prospects for observing such tumblers at the LHC, and we shall assess the extent to which such tumbler signatures can be distinguished from other signatures of new physics which also involve multiple DVs within the same collider event. We shall also develop a procedure for reconstructing the masses and lifetimes of the particles involved in the corresponding decay chains.

One important theme running through this work will be the observation that the prospects for discovering and distinguishing tumblers can be greatly enhanced by exploiting precision timing information. Fortunately, this sort of information can be provided by a precision timing layer of the sort that will be installed within the CMS detector during the forthcoming high-luminosity upgrade of the Large Hadron Collider (LHC) [6,7]. As we shall see, this timing information can significantly improve the precision with which the masses and lifetimes of the particles within a tumbler can be measured.

This paper is organized as follows. In Sec. II, we describe the basic properties of tumblers and discuss the role that timing information can play in characterizing them. In Sec. III, we introduce a concrete example model which can give rise to tumblers. In Sec. IV, we survey the parameter space of this model and identify regions of this parameter space wherein the prospects for identifying tumblers are particularly auspicious. In Sec. V, we investigate the extent to which current LHC data constrains this parameter space and assess the prospects for observing a significant number of tumbler events both before and after the high-luminosity LHC (HL-LHC) upgrade. In Sec. VI, we develop an event-selection procedure which provides an efficient way of distinguishing between events which involve tumblers and events which involve multiple DVs which were not in fact produced by the successive decays of unstable particles within the same decay chain. We also investigate the degree to which the masses and lifetimes of the dark-sector states can be measured from tumbler events. In Sec. VII, we conclude with a summary of our results and a discussion of the ways in which improvements in energy and timing resolution could enhance our ability to distinguish tumbler signatures at the HL-LHC or at future colliders.

II. TUMBLERS AT THE LHC

Macroscopically displaced vertices can result from the decays of long-lived particles (LLPs) that decay on distance scales $\mathcal{O}(1 \text{ mm}) \lesssim c\tau \lesssim \mathcal{O}(100 \text{ m})$ inside a collider detector. Such vertices represent a striking potential signal of new physics [1,4]. While the DV signatures associated with the decays of even a single LLP species can yield a wealth of information about physics beyond the SM, the phenomenology associated with DVs can be far richer in extensions of the SM which involve multiple of LLP species. One intriguing possibility arises in scenarios in which one species of LLP can decay into a final state which includes both SM particles and a lighter LLP of a different species. If this lighter LLP also decays within the detector, the result is a sequence of two or more DVs which result from the successive decays of unstable particles within the same decay chain. Like DVs themselves, such *sequences* of DVs—i.e., such “tumblers”—can arise naturally in many extensions of the SM. These include models involving compressed supersymmetry [8]; hidden-valley models [9] and other, similar theories which give rise to phenomena referred to as emerging jets [5], semivisible jets [10], dark jets [11], and/or soft bombs [12]; theories involving large numbers of additional degrees of freedom with a significant degree of disorder in their mass matrix [13]; and scenarios involving nonminimal dark sectors [14]. Indeed, tumbler-like events of this sort, under the name “microcascades,” were explicitly invoked more than a decade ago to explain possible anomalies in CDF data involving muons produced with large impact parameters [15,16]. The possibility of tumbler-like events arising in a variety of hidden-valley models was also discussed in Refs. [17–20].

An example of a tumbler is illustrated in Fig. 1. In this example, an LLP χ_2 is produced within a collider detector at the primary interaction vertex V_P , along with one or more additional SM particles. This χ_2 particle travels a measurable distance away from V_P before it decays into a pair of SM particles (which for concreteness we take to be a quark q and an antiquark \bar{q}), along with another, lighter LLP χ_1 at the secondary vertex V_S . This χ_1 particle, in turn, travels a measurable distance away from V_S before it likewise decays at a tertiary vertex V_T into a quark q' , an antiquark \bar{q}' , and another, even lighter LLP χ_0 , which escapes the detector and manifests itself as missing transverse energy \cancel{E}_T .

Figure 1 illustrates the topology of a tumbler involving only two DVs, as appropriate for a decay chain involving three LLPs (χ_2 , χ_1 , and χ_0). In some sense, this is the minimal possible tumbler, and this case will be the focus of this paper. However, there is nothing that requires tumblers to be limited to only two DVs or three LLPs, and indeed longer decay chains leading to more DVs are possible. Indeed, in many SM extensions, entire ensembles of LLPs χ_n can arise. Such ensembles can then give rise to potentially long decay chains with many sequential DVs.




FIG. 1. Schematic of a tumbler event within a collider detector modeled after the CMS detector at the HL-LHC. In this event, a heavy LLP χ_2 is initially produced at the primary vertex V_p , along with some additional particles. The χ_2 particle then travels a measurable distance before decaying into a lighter LLP χ_1 , a quark q , and an antiquark \bar{q} at the secondary vertex V_s . This χ_1 particle then travels a measurable distance away from V_s and subsequently decays into an even lighter LLP χ_0 , another quark q' , and another antiquark \bar{q}' at the tertiary vertex V_T . The χ_0 particle manifests itself as missing energy \cancel{E}_T , while the quarks and antiquarks manifest themselves as hadronic jets. Information about when each jet interacts with the timing layer, in conjunction with additional information about the momentum of the jet from the tracker and calorimeters, can be used to reconstruct the locations and times at which V_s and V_T occurred.

However, all such tumbler events share the same basic event topology, with sequential decays proceeding in linear fashion down the decay chain.

How might such a tumbler be detected and distinguished? Since the SM backgrounds for processes involving DVs are quite low, signals involving DVs provide particularly striking indications of new physics. A variety of LLP searches involving DVs have already been performed by the ATLAS and CMS Collaborations. Moreover, the sensitivity of the ATLAS and CMS detectors to DV signatures will be significantly enhanced during the forthcoming HL-LHC upgrade, in part as a result of the installation of additional apparatus within both of these detectors which provides precision timing information about the particles produced in a collider event. In particular, the upgraded ATLAS detector will include a high-granularity timing detector in front of each of the endcap calorimeters in order to provide timing information for particles emitted in the forward direction [21]. The upgraded CMS detector, by contrast, will include not only a pair of timing detectors located in front of the endcaps, but also a thin cylindrical timing layer situated between the tracker and the electromagnetic calorimeter (ECAL) which provides coverage within the barrel region of the detector [6,7]. This timing layer, which is included in the illustration in Fig. 1, will provide a timing resolution of $\sigma_t \approx 30$ ps—a vast improvement over the timing resolution $\sigma_t \approx 150$ ps currently afforded by the ECAL itself [22]. Such a

significant enhancement in timing precision will significantly improve the sensitivity of LLP searches at the HL-LHC. Indeed, not only can information from the timing layer be used to reduce SM backgrounds for such searches [23,24], but it can also aid in the reconstruction of the LLP masses [25], strategies for which have been developed for a number of LLP-decay scenarios [26,27]. In particular, the momenta \vec{p}_q and $\vec{p}_{\bar{q}}$ of the hadronic jets associated with q and \bar{q} , in conjunction with timing the information for these jets provided by either the timing layer or the ECAL, can be used to identify both the time t_s and spatial location \vec{x}_s of V_s . Similarly, the momenta $\vec{p}_{q'}$ and $\vec{p}_{\bar{q}'}$ of the jets associated with q' and \bar{q}' , in conjunction with the corresponding timing information, can be used to identify the time t_T and spatial location \vec{x}_T of V_T . Information about the momenta of the additional SM particles produced at V_p , in conjunction with the corresponding timing information, can be used to identify the time t_p and spatial position \vec{x}_p of this vertex.

III. A CONCRETE EXAMPLE MODEL

In order to perform a more quantitative assessment of the prospects for detecting tumbler signatures at the LHC and beyond, it is necessary to work within the context of a concrete model. Such a model can therefore also serve as an existence proof that tumblers may indeed arise at colliders such as the LHC, and yet be consistent with current

experimental results. The model that we adopt for this purpose is drawn from a general class of nonminimal dark-sector scenarios in which there exist multiple dark-sector states χ_n with similar quantum numbers, all of which can interact with the fields of the visible sector via a common mediator particle ϕ . Not only do these interactions provide a portal through which the χ_n can be produced, but they also render the heavier χ_n unstable. Since the final states into which these χ_n decay in such scenarios generically involve both SM particles and other, lighter dark-sector states χ_m , extended decay chains can develop.

In Ref. [14] we constructed such a model within this class and focused on a region of parameter space in which the χ_n particles involved in these decay chains had lifetimes leading to prompt decays rather than macroscopically displaced vertices. We then discovered that such decay chains can lead to striking signatures involving large multiplicities of produced SM states.

In this paper, by contrast, we shall focus on a different region within the parameter space of this model, one in which the χ_n have lifetimes within the range $\mathcal{O}(1 \text{ mm}) \lesssim c\tau_n \lesssim \mathcal{O}(100 \text{ m})$. As we shall explain further below, we thus obtain decay chains involving DVs—i.e., tumblers. Moreover, although our analysis in Ref. [14] considered arbitrary numbers of χ_n states within the associated decay chains, we shall here restrict our attention to cases with only three χ_n particles, with $n = 0, 1, 2$ labeling these states in order of increasing mass.

More specifically, this model is defined as follows. We shall take the χ_n to be Dirac fermions and to be singlets under the SM gauge group. We take the masses m_n of the χ_n to be free parameters, subject to the condition $m_2 > m_1 > m_0$. The particle ϕ which mediates the interactions between the χ_n and the fields of the SM in our model is taken to be a complex scalar which transforms as a triplet both under the SM $SU(3)_c$ gauge group and under the approximate $U(3)_u$ flavor symmetry of the right-handed up-type quarks. In order to alleviate issues involving flavor-changing effects, we shall assume that the up-type quarks $q \in \{u, c, t\}$ and the component fields ϕ_q within ϕ share a common mass eigenbasis. Expressed in this eigenbasis, the interaction Lagrangian which couples the dark and visible sectors then takes the form

$$\mathcal{L}_{\text{int}} = \sum_q \sum_{n=0}^2 [c_{nq} \phi_q^\dagger \bar{\chi}_n P_R q + \text{H.c.}], \quad (3.1)$$

where $P_R \equiv \frac{1}{2}(1 + \gamma^5)$ is the usual right-handed projection operator and where c_{nq} is a dimensionless coupling constant which in principle depends both on the value of the index n for the dark-sector field and on the flavor of the quark. Such a coupling structure implies that each of the ϕ_q couples only to a single quark flavor q .

For simplicity, we shall focus on the case in which the masses of the mediators ϕ_c and ϕ_t which couple to the charm and top quarks are sufficiently large that they greatly exceed the mass of the mediator ϕ_u (i.e., $m_{\phi_c}, m_{\phi_t} \gg m_{\phi_u}$) and also have no appreciable impact on the collider phenomenology of the model. From a low-energy perspective, this is equivalent to adopting a coupling structure in which $c_{nc} \approx 0$ and $c_{nt} \approx 0$ for all n , while the $c_n \equiv c_{nu}$ are in general nonvanishing. Moreover, for concreteness, we shall assume that the c_n scale according to the power-law relation

$$c_n = c_0 \left(\frac{m_n}{m_0} \right)^\gamma, \quad (3.2)$$

where c_0 is the coupling associated with the lightest ensemble constituent χ_0 and where γ is a dimensionless scaling exponent.

In summary, our model is characterized by six free parameters. These are the masses m_n of the three χ_n , the parameters c_0 and γ which specify the couplings between these fields and the mediator ϕ_u , and the mass m_{ϕ_u} of the mediator itself. For ease of notation, since we are assuming that ϕ_c and ϕ_t are sufficiently heavy that they play no role in the collider phenomenology of the model, we shall henceforth simply refer to ϕ_u and m_{ϕ_u} as ϕ and m_ϕ , respectively.

Our interest in this model is primarily due to the tumbler signatures which result from successive decays of the dark-sector states. Indeed, the interaction Lagrangian in Eq. (3.1) renders χ_1 and χ_2 unstable. We shall primarily be interested in the regime within which the mediator is heavy, with $m_\phi > m_n$ for all n . Within this regime, the leading contribution to the decay width Γ_ϕ of the mediator arises from to two-body decay processes of the form $\phi \rightarrow q\bar{\chi}_n$. By contrast, the leading contribution to the decay widths of each χ_1 and χ_2 arise from three-body decay processes of the form $\chi_n \rightarrow q\bar{q}\chi_m$ with $m < n$, each of which involves an off-shell mediator. Thus, when a χ_2 particle is produced at the primary interaction vertex V_P , there is a nonvanishing probability that it will decay via the process $\chi_2 \rightarrow q\bar{q}\chi_1$, with χ_1 in turn decaying via the process $\chi_1 \rightarrow q\bar{q}\chi_0$. The resulting decay chain is illustrated in Fig. 2, where each black dot represents an interaction vertex associated with




FIG. 2. Realization of the tumbler event topology shown in Fig. 1 within the context of our model. In particular, within our model, the secondary and tertiary vertices V_S and V_T in Fig. 1 are each now effectively realized as a pair of three-point vertices mediated by ϕ .

one of the Lagrangian terms in Eq. (3.1). Since the ϕ particles involved in the decay processes are both off shell, the red circles indicated in the diagram, each of which encompasses two such interaction vertices, represent localized spacetime events. If the χ_1 and χ_2 particles are both long-lived and each travel a macroscopic distance before they decay, the result is a tumbler, with these spacetime events corresponding to the secondary and tertiary decay vertices V_S and V_T indicated in Fig. 1.

Although χ_1 and χ_2 are unstable, the lightest dark-sector state χ_0 in our model is stabilized by an accidental \mathbb{Z}_2 symmetry of the model under which ϕ and the χ_n are odd, whereas the fields of the SM are even. This symmetry, if unbroken, would render this particle absolutely stable—and a potential dark-matter candidate [14]. Alternatively, this symmetry could be broken by additional, highly suppressed interactions which permit χ_0 to decay into final states involving SM particles alone. However, as long as χ_0 is collider-stable—i.e., sufficiently long-lived that virtually every χ_0 particle produced within a collider detector escapes the detector well before it decays—we shall not need to specify whether this \mathbb{Z}_2 symmetry is exact or approximate for the purposes of understanding the collider phenomenology of the model. In what follows, we shall therefore simply assume that χ_0 is indeed collider-stable and consequently manifests itself as \cancel{E}_T .

IV. SURVEYING THE PARAMETER SPACE

Our first step is to identify regions of the parameter space of our model within which the prospects for observing a tumbler signature, either at the LHC or at a future hadron collider, are particularly auspicious. The event rate for collider processes involving tumblers depends on several factors. These include the cross sections for the relevant production processes; the lifetimes of χ_1, χ_2 , and ϕ ; and the probability that an on-shell ϕ or χ_2 particle initially produced via one of these production processes will decay via an appropriate decay chain.

We begin by evaluating the total decay widths Γ_ϕ and Γ_n and the branching fractions $\text{BR}_{\phi n}$ and $\text{BR}_{n\ell}$ for decay processes of the form $\phi^\dagger \rightarrow \bar{q}\chi_n$ and $\chi_n \rightarrow \bar{q}q\chi_\ell$, respectively. In order to calculate these branching fractions, we must first evaluate the partial widths for all kinematically accessible decays of ϕ, χ_2 , and χ_1 . The partial width $\Gamma_{\phi n} \equiv \Gamma(\phi^\dagger \rightarrow \bar{q}\chi_n)$ for the decay process in which an on-shell mediator decays into a quark and an ensemble constituent χ_n is [14]

$$\Gamma_{\phi n} = \frac{c_n^2}{16\pi} \frac{(m_\phi^2 - m_n^2)^2}{m_\phi^3}. \quad (4.1)$$

Likewise, the partial width $\Gamma_{n\ell} \equiv \Gamma(\chi_n \rightarrow \bar{q}q\chi_\ell)$ takes the form [14]

$$\Gamma_{n\ell} = \frac{3c_n^2 c_\ell^2 m_\phi}{256\pi^2 r_{\phi n}^3} \left[f_{\phi n\ell}^{(1)} - f_{\phi n\ell}^{(2)} \ln(r_{n\ell}) + f_{\phi n\ell}^{(3)} \ln\left(\frac{1 - r_{\phi n}^2}{1 - r_{\phi n}^2 r_{n\ell}^2}\right) \right], \quad (4.2)$$

where $r_{n\ell} \equiv m_\ell/m_n$, where $r_{\phi n} \equiv m_n/m_\phi$, and where

$$\begin{aligned} f_{\phi n\ell}^{(1)} &\equiv 6r_{\phi n}^2(1 - r_{n\ell}^2) - 5r_{\phi n}^4(1 - r_{n\ell}^4) + 2r_{\phi n}^6 r_{n\ell}^2(1 - r_{n\ell}^2), \\ f_{\phi n\ell}^{(2)} &\equiv 4r_{\phi n}^8 r_{n\ell}^4, \\ f_{\phi n\ell}^{(3)} &\equiv 6 - 8r_{\phi n}^2(1 + r_{n\ell}^2) - 2r_{\phi n}^8 r_{n\ell}^4 + 2r_{\phi n}^4(1 + 4r_{n\ell}^2 + r_{n\ell}^4). \end{aligned} \quad (4.3)$$

The branching fractions of interest are then given by

$$\text{BR}_{\phi n} = \frac{\Gamma_{\phi n}}{\Gamma_\phi}, \quad \text{BR}_{n\ell} = \frac{\Gamma_{n\ell}}{\Gamma_n}, \quad (4.4)$$

where the total widths of ϕ and χ_n are respectively given by

$$\Gamma_\phi = \sum_{n=0}^2 \Gamma_{n\phi}, \quad \Gamma_n = \sum_{\ell=0}^{n-1} \Gamma_{n\ell}. \quad (4.5)$$

We observe from the partial-width expressions in Eqs. (4.1) and (4.2) that $\Gamma_\phi \propto c_0^2$, while $\Gamma_n \propto c_0^4$. For $m_n \sim \mathcal{O}(100 \text{ GeV})$ and $m_\phi \sim \mathcal{O}(\text{TeV})$, these expressions also imply that we must take $c_0 \ll 1$ in order for χ_1 and χ_2 to be sufficiently long-lived that their decays give rise to DVs. Together, these two considerations imply that $\Gamma_\phi \gg \Gamma_n$ within regions of parameter space which give rise to tumblers. As a result, within these regions of interest, any on-shell ϕ particle produced at the primary interaction vertex typically decays promptly into a quark and one of the χ_n .

From the branching fractions in Eq. (4.4), we may in turn determine the probability that a particular decay chain will arise from the decay of a ϕ or χ_n particle. We shall let $P_{a_1 a_2 \dots a_f}$ denote the probability of a given decay chain, where the sequence of $a_i \in \{\phi, 2, 1, 0\}$ in the subscript indicates the set of ϕ and χ_n particles produced along the decay chain. For example, $P_{\phi 20}$ represents the probability that an on-shell ϕ particle, once produced, decays directly to χ_2 , which subsequently decays directly to χ_0 . These decay-chain probabilities are simply the products of the relevant branching fractions. Since χ_1 decays via the process $\chi_1 \rightarrow \bar{q}q\chi_0$ with branching fraction $\text{BR}_{10} = 1$, we have $P_{10} = 1$. There are two possible decay chains which can arise from the decay of a χ_2 particle, given that χ_2 can decay either to χ_0 directly, or to χ_1 which then subsequently decays to χ_0 . The respective decay-chain probabilities are therefore $P_{20} = \text{BR}_{20}$ and $P_{210} = \text{BR}_{21}$. The probabilities associated with decay chains initiated by

the decays of ϕ and χ_2 can be evaluated in a similar manner. We emphasize that each of these decay-chain probabilities represents the *total* probability associated with the corresponding sequence of decays *regardless* of the likelihood that these decays would occur within a collider detector.

We now consider the production processes through which ϕ and χ_n particles can be produced at a hadron collider. The accidental \mathbb{Z}_2 symmetry of our model ensures that particles which are odd under this symmetry will always be produced in pairs. The dominant scattering processes which give rise to a signal in our toy model are therefore $pp \rightarrow \phi^\dagger\phi$, $pp \rightarrow \phi\chi_n$ (and its Hermitian-conjugate process), and $pp \rightarrow \bar{\chi}_m\chi_n$. The Feynman diagrams which provide the leading contributions to the cross sections for these processes are shown in Ref. [14].

Since ϕ carries color charge, the dominant contribution to the cross-section $\sigma_{\phi\phi}$ for the process $pp \rightarrow \phi^\dagger\phi$ comes from diagrams which involve strong interactions alone. By contrast, the diagrams which provide the dominant contribution to the cross-section $\sigma_{\phi n}$ for any process of the form $pp \rightarrow \phi\chi_n$ each include one vertex which follows from the interaction Lagrangian in Eq. (3.1). Likewise, the diagrams which provide the dominant contribution to the cross-section σ_{mn} for any process of the form $pp \rightarrow \bar{\chi}_m\chi_n$ each include two such vertices. These considerations imply that $\sigma_{\phi\phi}$ is independent of c_0 , while $\sigma_{\phi n} \propto c_0^2$ and $\sigma_{mn} \propto c_0^4$. Thus, since $c_0 \ll 1$ within regions of parameter space which give rise to tumblers, $pp \rightarrow \phi^\dagger\phi$ typically dominates the production rate for tumbler events by several orders of magnitude within those regions.¹ As a result, while the branching fractions $\text{BR}_{\phi n}$ and $\text{BR}_{n\ell}$ depend on the values of γ , c_0 , m_0 , m_1 , and m_2 , the cross-section $\sigma_{\phi\phi}$ for the sole scattering process relevant for tumbler production at hadron colliders depends essentially on m_ϕ alone.

Since $pp \rightarrow \phi^\dagger\phi$ typically provides the dominant contribution to the tumbler event rate within our parameter-space region of interest, it is the decays of on-shell mediator particles which typically provide the dominant contribution to the tumbler-event rate. The sole decay chain through which an on-shell ϕ particle, once produced by this process, can give rise to a tumbler is the chain in which this ϕ particle decays promptly to a χ_2 particle, which then decays to a χ_1 particle (which itself subsequently decays to a χ_0 particle with $\text{BR}_{10} = 1$). Thus, the decay-chain probability $P_{\phi 210} = \text{BR}_{\phi 2}\text{BR}_{21}$ for this sequence of decays is a crucial figure of merit in assessing whether or not a given choice of our model parameters is likely to lead to a significant number of tumbler events at a hadron collider.

¹In unusual circumstances wherein $\text{BR}_{\phi 2}$ is suppressed by phase-space considerations and ϕ decays do not tend to produce tumblers, it is also possible that $pp \rightarrow \phi\chi_2$ dominates this event rate. However, since this possibility requires that the masses m_2 and m_ϕ be tuned such that they are nearly equal, we do not consider it further.

In order to assess which regions of the parameter space of our model are the most promising for tumbler detection, we search for points at which the following criteria are satisfied. First, the proper decay distances $c\tau_1$ and $c\tau_2$ of the unstable LLPs must each lie within the range $1 \text{ mm} < c\tau_n < 10 \text{ m}$. These conditions ensure not only that a χ_1 or χ_2 particle has a significant probability of traveling an appreciable distance away from the location at which it was produced before it decays, but also that it has a significant probability of decaying before it leaves the detector tracker. Second, we require that $m_2 < m_\phi$ in order to ensure that the decay $\phi^\dagger \rightarrow \bar{q}\chi_2$ is kinematically allowed. Third, we require that $P_{\phi 210}$ exceed a certain threshold. In general, P can be as high as $P_{\phi 210} \sim \mathcal{O}(0.1)$; indeed, this occurs despite the fact that $P_{\phi 210}$ is often suppressed by phase-space considerations which favor the decay of ϕ , χ_2 , and χ_1 directly to χ_0 . That said, we shall nevertheless adopt the far more modest requirement $P_{\phi 210} \gtrsim 10^{-6}$ in our survey in order that we may better explore how this decay-chain probability varies across the parameter space as a whole.

In Fig. 3, we plot contours of $P_{\phi 210}$ in $(\Delta m_{10}, \Delta m_{21})$ space, where $\Delta m_{10} \equiv m_1 - m_0$ and $\Delta m_{21} \equiv m_2 - m_1$. Results are only shown for regions wherein all of the three criteria discussed above are satisfied; other regions appear in white. The results shown in the left panel correspond to the parameter assignments $m_\phi = 1750 \text{ GeV}$, $m_0 = 600 \text{ GeV}$, $c_0 = 0.001$, and $\gamma = 0$. The results shown in the right panel correspond to the same assignments for m_ϕ , m_0 , and c_0 , but with $\gamma = 1$.

Broadly speaking, within these regions, the largest values of $P_{\phi 210}$ are obtained when Δm_{10} is small and Δm_{21} is large. Moreover, we see that tumbler decay-chain probabilities as large as $P_{\phi 210} \sim \mathcal{O}(0.1)$ can arise within this region for $\gamma = 1$, whereas probabilities as large as $P_{\phi 210} \sim \mathcal{O}(0.01)$ can arise even for $\gamma = 0$. Within the white region on the left side of each panel, the available phase space for the decay $\chi_1 \rightarrow \bar{q}q\chi_0$ is extremely small, and consequently $c\tau_1 > 10 \text{ m}$. By contrast, within the white region in the upper right corner of each panel, m_2 is quite large. As a result, either the partial width for the decay $\chi_2 \rightarrow \bar{q}q\chi_0$ becomes so large that $c\tau_2 < 1 \text{ mm}$, or else $m_2 > m_\phi$ and the three-step decay chain which gives rise to tumblers is kinematically forbidden. While the results shown in Fig. 3 by no means represent an exhaustive survey of the parameter space of our model, they serve to highlight those regions which could potentially yield a significant number of tumbler events at the LHC or at future colliders.

Guided by these results, then, we shall identify a set of four benchmark points within these regions for further study. The parameter assignments which define these benchmark points are provided in Table I. Each point is also labeled with a star in Fig. 3. These benchmark points




FIG. 3. Contours within the $(\Delta m_{10}, \Delta m_{21})$ -plane of the overall probability $P_{\phi 210} = \text{BR}_{\phi 2} \text{BR}_{21}$ that an on-shell mediator ϕ will decay via the three-step decay chain which yields a tumbler. The results shown in the left panel correspond to the parameter assignments $m_\phi = 1750$ GeV, $m_0 = 600$ GeV, $c_0 = 0.001$, and $\gamma = 0$. The results shown in the right panel correspond to the same assignments for m_ϕ , m_0 , and c_0 , but with $\gamma = 1$. Regions of parameter space shown in white are not of interest from a tumbler perspective, either because one of the relevant decay processes is kinematically forbidden, because one or both of the proper decay lengths $c\tau_1$ and $c\tau_2$ of the unstable LLPs lies below 1 mm or above 10 m, or because $P_{\phi 210} < 10^{-6}$. The four stars which appear in the panels of this figure indicate the parameter-space benchmarks defined in Table I.

represent different combinations of the parameters γ , m_1 , and m_2 . The mass splittings Δm_{10} and Δm_{21} , the branching fractions for the different possible decay channels for ϕ and χ_2 , and the proper decay lengths of χ_1 and χ_2 for each of these benchmarks are provided in Table II.

TABLE I. Definitions of our parameter-space benchmarks BM1–BM4.

Benchmark	Input parameters					
	c_0	γ	m_0 (GeV)	m_1 (GeV)	m_2 (GeV)	m_ϕ (GeV)
BM1	0.001	0	600	800	1000	1750
BM2	0.001	1	600	800	1000	1750
BM3	0.001	1	600	800	1200	1750
BM4	0.001	1	600	1000	1200	1750

It is also interesting to consider how our results for $P_{\phi 210}$ vary as a function of the choice of the scaling exponent γ . In Fig. 4, we plot contours of $P_{\phi 210}$ within the $(\Delta m_{10}, \gamma)$ -plane for $\Delta m_{21} = 200$ GeV (left panel) and $\Delta m_{21} = 400$ GeV (right panel). The values we have adopted for m_ϕ , m_0 , and c_0 in both panels of the figure are the same as those adopted in Fig. 3. The locations of our parameter-space benchmarks are once again indicated by the stars. We see that increasing γ with all other parameters held fixed generally increases $P_{\phi 210}$. Indeed, increasing this scaling exponent increases the ratios c_2/c_1 and c_2/c_0 , and thereby increases the branching fraction $\text{BR}_{\phi 2}$ for the decay $\phi \rightarrow q\bar{q}\chi_2$ that initiates the three-step decay chain which gives rise to tumblers. By the same token, however, increasing γ also increases the total decay width of χ_2 . For sufficiently large γ , the lifetime of this particle becomes such that $c\tau_2 < 1$ mm. This is what occurs in the white region in the

TABLE II. Values for the mass splittings $\Delta m_{10} \equiv m_1 - m_0$ and $\Delta m_{21} \equiv m_2 - m_1$, the branching fractions for all of the processes via which ϕ and χ_2 can decay, and the proper decay lengths $c\tau_1$ and $c\tau_2$ of the unstable LLPs for each of the parameter-space benchmarks defined in Table I.

Benchmark	Mass splittings		Branching fractions					Proper decay lengths	
	Δm_{10} (GeV)	Δm_{21} (GeV)	$\text{BR}_{\phi 2}$	$\text{BR}_{\phi 1}$	$\text{BR}_{\phi 0}$	BR_{21}	BR_{20}	$c\tau_2$ (m)	$c\tau_1$ (m)
BM1	200	200	0.24	0.34	0.42	0.05	0.95	8.33×10^{-2}	2.42
BM2	200	200	0.40	0.35	0.25	0.08	0.92	2.89×10^{-2}	1.36
BM3	200	400	0.37	0.37	0.26	0.28	0.72	2.14×10^{-3}	1.36
BM4	400	200	0.36	0.40	0.25	0.03	0.97	2.89×10^{-3}	3.15×10^{-2}




FIG. 4. Same as in Fig. 3, except that the contours of $P_{\phi 210}$ are shown within the $(\Delta m_{10}, \gamma)$ -plane for $\Delta m_{21} = 200$ GeV (left panel) and $\Delta m_{21} = 400$ GeV (right panel).

upper right corner of each panel. On the other hand, when $\gamma < 0$, the decay $\phi \rightarrow q\bar{q}\chi_0$ dominates the width of ϕ . As a result, $P_{\phi 210}$ decreases rapidly with γ until it drops below the threshold $P_{\phi 210} > 10^{-6}$, leading to the white region in the lower right of the plot. As in Fig. 3, the white region on the left side of each panel corresponds to the region in which the available phase space for the decay $\chi_1 \rightarrow \bar{q}q\chi_0$ is small and $c\tau_1 > 10$ m.

V. CONSTRAINTS AND EVENT RATES

In the previous section, we identified the parameter-space regions of our model which are particularly auspicious for producing tumblers. In this section, we focus on these parameter-space regions of interest and assess whether a substantial population of tumbler events could yet await us at the LHC, given that no significant excess in discovery channels involving multiple DVs has been observed to date.

One important consideration is that our model not only gives rise to tumblers, but also yields contributions to the event rates in several additional detection channels for new physics. These channels include the monojet + \cancel{E}_T channel, the multijet + \cancel{E}_T channel, and various channels involving displaced hadronic jets. The results of new-physics searches which have been performed in these channels by the ATLAS and CMS Collaborations place additional constraints on the parameter space of our model. Thus, we begin our analysis with a summary of the relevant constraints from these searches.

A. Displaced-vertex search constraints

A variety of searches for signatures of new physics involving displaced hadronic jets have been performed by both the CMS and ATLAS Collaborations. The CMS Collaboration, for example, has recently performed one search for displaced jets with 137 fb^{-1} of integrated

luminosity which incorporates timing information from the ECAL [28], as well as another, similar search with 132 fb^{-1} of integrated luminosity in which dedicated displaced-jet triggers and background-reduction techniques were applied [29]. A CMS search for displaced jets emanating from a pair of DVs resulting from the decays of pair-produced of LLPs was also recently performed with 140 fb^{-1} of integrated luminosity [30]. The results of these searches collectively supersede those from similar CMS searches for displaced jets performed at 36 fb^{-1} [31] and 38.5 fb^{-1} [32] of integrated luminosity. The extent to which machine-learning techniques could be used in order to further improve the reach of searches involving displaced jets was investigated in Ref. [33].

The ATLAS Collaboration has likewise performed a number of different searches for LLPs decaying into displaced jets. These include searches for events in which the decay which produces the jets occurs within the tracker [34], within the calorimeter [35], or in the muon chamber [36]. An ATLAS search has also been performed for multiple LLPs decaying to jets in the same event, where one LLP decays within the tracker and the other decays within the muon chamber [37]. All of these searches are performed with roughly 35 fb^{-1} of integrated luminosity, though the precise value of the integrated luminosity varies slightly among these searches. Owing primarily to the substantially lower integrated luminosity, these ATLAS searches are not as constraining as the CMS searches. For this reason, we focus on the results of the CMS searches in what follows.

The results in Refs. [28–30] collectively constrain new-physics scenarios involving LLPs with lifetimes τ_χ in the range $10^{-4} \text{ m} \lesssim c\tau_\chi \lesssim 10 \text{ m}$ which decay into final states involving hadronic jets. In particular, they impose an upper bound on the product $\sigma_{\chi\chi} \text{BR}_{\chi,j}^2$ of the LLP pair-production cross section and the square of the branching fraction of the LLP into such final states. While the precise numerical

value of this upper bound depends on the production and decay kinematics of the LLP and on τ_χ , the bound falls within the range 0.05–0.5 fb across almost this entire range of τ_χ .

B. Multijet search constraints

Searches performed by both the CMS and ATLAS Collaborations also place constraints on beyond-the-Standard-Model (BSM) contributions to the event rate for processes involving multiple hadronic jets and \cancel{E}_T . The searches most relevant for constraining the parameter space of our model are those designed to uncover evidence of heavy decaying particles—e.g., squarks and gluinos in supersymmetry. The leading CMS constraints from multijet + \cancel{E}_T searches are those derived from searches [38,39] performed with 137 fb^{-1} of integrated luminosity. These include searches involving standard techniques developed in order to search for squarks and gluinos more generally, as well as searches which focus on specific scenarios for which the use of the M_{T2} variable is particularly advantageous in terms of discovery potential. The results of these analyses supersede those of a prior CMS study [40] performed with 36 fb^{-1} of integrated luminosity.

The leading ATLAS constraints on excesses in the multijet + \cancel{E}_T channel of the sort obtained in our model are those derived from a search for squarks and gluinos performed with 139 fb^{-1} of integrated luminosity [41]. These results supersede those obtained from a prior ATLAS study [42] performed with 36 fb^{-1} of integrated luminosity.

In each of these ATLAS or CMS analyses, 95%-C.L. exclusion limits on the product of the production cross-section σ , the signal acceptance A , and the detection efficiency ϵ are obtained for a variety of signal regions, which are defined differently in the different studies. These limits are also interpreted in each case as constraints on the parameter space of a simplified supersymmetric model involving a single flavor of squark \tilde{q} which is pair-produced via the process $pp \rightarrow \tilde{q}^\dagger \tilde{q}$ and subsequently decays directly to a light quark and the lightest neutralino $\tilde{\chi}_1$. All other sparticles are assumed to be extremely heavy in this scenario, and therefore to play no role in the pair-production process. Since \tilde{q} and $\tilde{\chi}_1$ in this supersymmetric model have the same quantum numbers as ϕ and χ_0 in our model, respectively, these bounds may be applied to our model directly. The constraint contours within the $(m_{\tilde{q}}, m_{\tilde{\chi}_1})$ -plane obtained in Refs. [38,39,41] are all roughly commensurate and, roughly speaking, exclude the region of this plane wherein $m_{\tilde{q}} \lesssim 1250 \text{ GeV}$ and $m_{\tilde{\chi}_1} \lesssim 500 \text{ GeV}$.

Given that the values of the parameters m_ϕ and m_0 for all of our parameter-space benchmarks lie well outside the corresponding region in the (m_ϕ, m_0) -plane, we may safely assume that our benchmarks are consistent with these constraints. Moreover, in many of these searches, events are vetoed in which a significant fraction of the jets are produced at locations other than the primary vertex.

C. Monojet search constraints

The most stringent bound on excesses of events in the monojet + \cancel{E}_T channel is that from an ATLAS study [43] performed with 139 fb^{-1} of integrated luminosity. The results of this study supersede those from a similar ATLAS study [44] performed with 36 fb^{-1} of integrated luminosity. Similar searches have been performed by the CMS Collaboration, but with far lower integrated luminosity.

The results in Ref. [43] are quoted in a model-independent way for several different signal regions corresponding with different threshold values taken for the magnitude $|\vec{p}_T^{(\text{rec})}|$ of the transverse momentum which recoils against the jet. For each of these signal regions, a 95%-C.L. exclusion limit on the product of the production cross-section σ , signal acceptance A , and detection efficiency ϵ is obtained. These limits range from $\sigma \times A \times \epsilon < 736 \text{ fb}$ for a threshold of $|\vec{p}_T^{(\text{rec})}| > 200 \text{ GeV}$ to $\sigma \times A \times \epsilon < 0.3 \text{ fb}$ for $|\vec{p}_T^{(\text{rec})}| > 1200 \text{ GeV}$. Moreover, these limits are also interpreted as constraints on the parameter space of the same simplified supersymmetric model that was considered in the multijet analysis discussed above. Once again, these constraints may be applied to our model directly.

The monojet constraints on this simplified supersymmetric model turn out to be relevant within the same rough region of the $(m_{\tilde{q}}, m_{\tilde{\chi}_1})$ -plane as the multijet constraints discussed above, but also are slightly less restrictive. We therefore expect that the same is true of the monojet constraints on our example model within the (m_ϕ, m_0) -plane. Thus, we may safely assume that our benchmarks are consistent with these constraints. In summary, then, it is clear that the dominant constraints on our model within our parameter-space region of interest are those from displaced-jet searches. We shall therefore focus primarily on these constraints in what follows.

D. Effective cross sections and event rates

In order to assess the impact of these experimental constraints on our model, we must evaluate the net contributions to the event rates for a number of different detection channels. In particular, we can identify four relevant channels, each of which is associated with a particular set of collider processes:

- Tumbler class: processes which involve at least one tumbler. Processes in this class are the primary focus of this paper.
- DV class: processes which involve at least one DV, regardless of whether this DV is part of a tumbler. The event rates associated with processes in this class are constrained by the results of displaced-jet searches.
- Multijet class: processes which do not give rise to any DVs, but instead yield a pair of prompt hadronic jets and missing transverse energy. Processes in this class contribute to the event rate in the multijet + \cancel{E}_T channel.
- Monojet class: processes which do not give rise to any DVs, but instead yield a single prompt hadronic jet and

TABLE III. List of the possible event topologies which can arise within our model from pair-production process of the form $pp \rightarrow \phi\phi$, $pp \rightarrow \phi\chi_n$, and $pp \rightarrow \chi_m\chi_n$. The entries in each column describe the corresponding properties of these topologies, with notation as described in the text.

First chain	Second chain	Tumblers	Displaced vertices	Prompt jets
from $pp \rightarrow \phi\phi$ production				
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	2T		2j
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\phi \rightarrow \chi_2 \rightarrow \chi_0$	T	DV	2j
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\phi \rightarrow \chi_1 \rightarrow \chi_0$	T	DV	2j
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\phi \rightarrow \chi_0$	T		2j
$\phi \rightarrow \chi_2 \rightarrow \chi_0$	$\phi \rightarrow \chi_2 \rightarrow \chi_0$		2DV	2j
$\phi \rightarrow \chi_2 \rightarrow \chi_0$	$\phi \rightarrow \chi_1 \rightarrow \chi_0$		2DV	2j
$\phi \rightarrow \chi_2 \rightarrow \chi_0$	$\phi \rightarrow \chi_0$		DV	2j
$\phi \rightarrow \chi_1 \rightarrow \chi_0$	$\phi \rightarrow \chi_2 \rightarrow \chi_0$		2DV	2j
$\phi \rightarrow \chi_1 \rightarrow \chi_0$	$\phi \rightarrow \chi_1 \rightarrow \chi_0$		2DV	2j
$\phi \rightarrow \chi_0$	$\phi \rightarrow \chi_0$			2j
from $pp \rightarrow \phi\chi_n$ production				
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	2T		j
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_0$	T	DV	j
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_1 \rightarrow \chi_0$	T	DV	j
$\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	χ_0	T		j
$\phi \rightarrow \chi_2 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	T	DV	j
$\phi \rightarrow \chi_2 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_0$		2DV	j
$\phi \rightarrow \chi_2 \rightarrow \chi_0$	$\chi_1 \rightarrow \chi_0$		2DV	j
$\phi \rightarrow \chi_2 \rightarrow \chi_0$	χ_0		DV	j
$\phi \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	T	DV	j
$\phi \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_0$		2DV	j
$\phi \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_1 \rightarrow \chi_0$		2DV	j
$\phi \rightarrow \chi_1 \rightarrow \chi_0$	χ_0		DV	j
$\phi \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	T		j
$\phi \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_0$		DV	j
$\phi \rightarrow \chi_0$	$\chi_1 \rightarrow \chi_0$		DV	j
$\phi \rightarrow \chi_0$	χ_0			j
from $pp \rightarrow \chi_m\chi_n$ production				
$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	2T		
$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_0$	T	DV	
$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	$\chi_1 \rightarrow \chi_0$	T	DV	
$\chi_2 \rightarrow \chi_1 \rightarrow \chi_0$	χ_0	T		
$\chi_2 \rightarrow \chi_0$	$\chi_2 \rightarrow \chi_0$		2DV	
$\chi_2 \rightarrow \chi_0$	$\chi_1 \rightarrow \chi_0$		2DV	
$\chi_2 \rightarrow \chi_0$	χ_0		DV	
$\chi_1 \rightarrow \chi_0$	$\chi_1 \rightarrow \chi_0$		2DV	
$\chi_1 \rightarrow \chi_0$	χ_0		DV	
χ_0	χ_0			

missing transverse energy. Processes in this class contribute to the event rate in the monojet + E_T channel. We emphasize that these classes are not mutually exclusive. For example, all processes in the tumbler class necessarily include DVs and are therefore also part of the DV class. We also emphasize that all processes within a particular class are not completely equivalent. One example of this is that the contributions from some DV-class processes may not be as stringently constrained by existing DV searches as the contributions from other such processes as a consequence of differences in kinematics and the event-selection criteria

involved. Another example, as we shall discuss further in Sec. VI, is that tumbler-class processes in which one or more additional hard jets are produced at the primary vertex are significantly more useful for reconstructing the masses and lifetimes of the χ_n . Nevertheless, as we shall see, this classification is useful in categorizing the contributions from our model to the event rates in different detection channels.

Contributions to the total event rate for each of these four classes of processes can in principle arise from a variety of different event topologies—i.e., different combinations of production processes. In Table III, we list all possible such

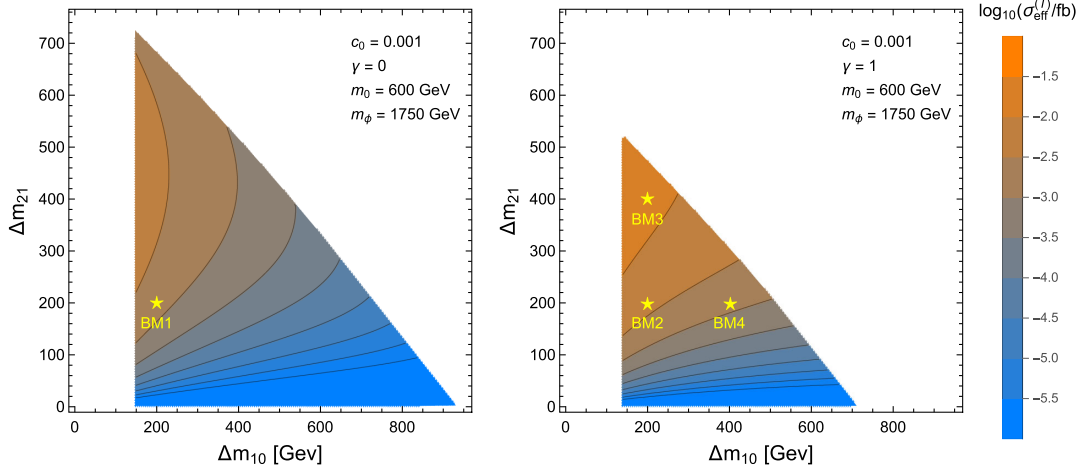


FIG. 5. Contours within the $(\Delta m_{10}, \Delta m_{21})$ -plane of the effective cross-section $\sigma_{\text{eff}}^{(T)}$ defined in Eq. (5.1) for processes involving at least one tumbler at the $\sqrt{s} = 14$ TeV HL-LHC. The results displayed in the left and right panels correspond to the parameter assignments in corresponding panels of Fig. 3. As in Fig. 3, results are shown only within regions wherein all decay processes involved in the production of a tumbler are kinematically allowed, where $c\tau_1$ and $c\tau_2$ both satisfy the criterion $1 \text{ mm} < c\tau_n < 10 \text{ m}$, and where $c\tau_\phi < 0.1 \text{ mm}$. The four stars indicate the locations of the parameter-space benchmarks defined in Table I.

event topologies which can arise from pair-production processes of the forms $pp \rightarrow \phi\phi$, $pp \rightarrow \phi\chi_n$, and $pp \rightarrow \chi_m\chi_n$. The first column indicates the structure of the longer decay chain in the event, while the second column indicates the structure of the shorter decay chain. An additional jet is produced by the decay of each mediator, while an additional pair of jets is produced by the decay of each LLP. However, for clarity, we have omitted mention of these particles in these columns of the table. Moreover, since there is no heuristic difference in terms of collider phenomenology between the decay chains precipitated by the decays of ϕ and χ_n and the decay chains precipitated by the decays of their antiparticles ϕ^\dagger and $\bar{\chi}_n$, we do not distinguish between particle and antiparticle decay chains. The third column of the table indicates whether the process gives rise to one or more tumblers at a collider. An entry of “T” in this column indicates that the process gives rise to a single tumbler, while an entry of “2T” indicates that the process gives rise to two tumblers, one from each decay chain. Likewise, the fourth column indicates whether or not the process gives rise to an isolated DV—i.e., a DV which is *not* part of a tumbler. An entry of “DV” in this column indicates the presence of a single such vertex, while an entry of “2DV” indicates the presence of such vertices. Finally, the fifth column indicates the presence of one or more prompt jets in the event. An entry of “j” indicates the presence of one such jet, while an entry of “2j” indicates the presence of two such jets. We note that since every decay chain which occurs in our model terminates with χ_0 , every event which results from any of the processes listed in this table also includes \cancel{E}_T .

For each of the four class of processes α itemized above, we define an effective cross-section $\sigma_{\text{eff}}^{(\alpha)}$ which represents

the sum of the individual contributions from all combinations of production and decay processes listed in Table III that contribute to the overall event rate for processes in that class. Each such individual contribution to $\sigma_{\text{eff}}^{(\alpha)}$ is the product of the cross-section $\sigma_{a_1 a_2}$ for the pair-production process $pp \rightarrow a_1 a_2$, where $a_i \in \{\phi, 2, 1, 0\}$, and the two decay-chain probabilities P_{a_1, c_1} and P_{a_2, c_2} associated with the decay chains on each side of the event. The index c_i appearing in these probabilities represents the sequence of particles produced from the decay of the corresponding initial particle a_i and includes the null decay chain in the event that the initial particle is stable, in which case the corresponding decay-chain probability is unity. In other words, our effective cross section is

$$\sigma_{\text{eff}}^{(\alpha)} \equiv \sum_{a_1} \sum_{a_2} \sum_{c_1} \sum_{c_2} [\sigma_{a_1 a_2} P_{a_1, c_1} P_{a_2, c_2}]_{\alpha}, \quad (5.1)$$

where the subscript α on the brackets enclosing the summand indicates that only event topologies associated with the corresponding class of processes are included in the sum. Indeed, it is the product of this effective cross section and the integrated luminosity which yield the overall event count for the corresponding class of processes.

In Fig. 5, we show contours of the effective cross-section $\sigma_{\text{eff}}^{(T)}$ for tumbler-class processes in $(\Delta m_{10}, \Delta m_{21})$ -space. Cross sections for all of the individual production processes were computed using the MG5_aMC@NLO code package [45] for a center-of-mass energy $\sqrt{s} = 14$ TeV. The results displayed in the left and right panels of the figure correspond to the parameter assignments in the corresponding panels of Fig. 3. As in Fig. 3, results are shown only

TABLE IV. The effective cross-sections $\sigma_{\text{eff}}^{(\alpha)}$ for tumbler-class, DV-class, and multijet-class processes for our parameter-space benchmarks. Also shown are the total numbers of tumbler events expected after Run 2 of the LHC and after the full HL-LHC run. We quote this number of events as $2\sigma_{\text{eff}}^{(\text{T})}\mathcal{L}_{\text{int}}$ in order to account for the contributions from both the CMS and ATLAS detectors.

Benchmark	$\sigma_{\text{eff}}^{(\alpha)}$ (fb)			Tumbler events	
	Tumblers	DV	Multijet + \cancel{E}_T	LHC Run 2 (137 fb $^{-1}$)	HL-LHC (3000 fb $^{-1}$)
BM1	1.5×10^{-3}	5.3×10^{-2}	1.1×10^{-2}	0.4	9.2
BM2	4.3×10^{-3}	6.1×10^{-2}	4.0×10^{-3}	1.1	25.6
BM3	1.3×10^{-2}	6.0×10^{-2}	4.3×10^{-3}	3.7	76.1
BM4	1.4×10^{-3}	6.1×10^{-2}	3.9×10^{-3}	0.4	8.1

within regions wherein all decay processes involved in the production of a tumbler are kinematically allowed, where the proper decay lengths $c\tau_1$ and $c\tau_2$ of the unstable LLPs both satisfy the criterion $1 \text{ mm} < c\tau_n < 10 \text{ m}$, and where the proper decay length of the mediator satisfies $c\tau_\phi < 0.1 \text{ mm}$. However, no minimum threshold for $P_{\phi 210}$ is imposed. The four stars once again indicate the locations of the parameter-space benchmarks defined in Table I.

We observe that the contours of $\sigma_{\text{eff}}^{(\text{T})}$ displayed in Fig. 5 have roughly the same shape as the contours of $P_{\phi 210}$ displayed in Fig. 3. This follows from the fact that $pp \rightarrow \phi^\dagger \phi$ vastly dominates the event rate within our parameter-space region of interest. As discussed in Sec. IV, the cross section for this process depends essentially on m_ϕ alone, and is therefore roughly uniform across the $(\Delta m_{10}, \Delta m_{21})$ -plane shown in each panel. More importantly, however, we also observe that an effective cross section of order $\sigma_{\text{eff}}^{(\text{T})} \sim \mathcal{O}(1\text{--}100 \text{ ab})$ for tumbler-class processes can be achieved across a substantial region of our parameters space—a region which includes the locations of all four of our parameter-space benchmarks. Given the integrated luminosity $\mathcal{L}_{\text{int}} = 3000 \text{ fb}^{-1}$ anticipated for the full HL-LHC run, cross sections of this order are in principle expected to give rise to a significant number of tumbler events at the HL-LHC.

In Table IV, we list the values of $\sigma_{\text{eff}}^{(\text{T})}$ obtained for each of our four benchmarks, along with the respective effective cross-sections $\sigma_{\text{eff}}^{(\text{DV})}$ and $\sigma_{\text{eff}}^{(\text{Nj})}$ for DV-class and multijet-class processes. Also shown in the figure are the corresponding total numbers of tumbler events expected after Run 2 of the LHC ($\mathcal{L}_{\text{int}} = 137 \text{ fb}^{-1}$) and after the full HL-LHC run ($\mathcal{L}_{\text{int}} = 3000 \text{ fb}^{-1}$). We quote this number of events as $2\sigma_{\text{eff}}^{(\text{T})}\mathcal{L}_{\text{int}}$ in order to account for the contributions from both the CMS and ATLAS detectors. While $\sigma_{\text{eff}}^{(\text{T})}$ varies significantly across the $(\Delta m_{10}, \Delta m_{21})$ -plane shown in the panels of Fig. 5, we find that $\sigma_{\text{eff}}^{(\text{DV})}$ and $\sigma_{\text{eff}}^{(\text{Nj})}$ are far less sensitive to the values of Δm_{10} and Δm_{21} within these same regions. Indeed, we find that both of these effective cross sections remain roughly within a single order of magnitude across this same region of $(\Delta m_{10}, \Delta m_{21})$ space.

One of the primary messages of Table IV is that the effective cross-section $\sigma_{\text{eff}}^{(\text{DV})}$ for each of our parameter-space benchmarks is $\sigma_{\text{eff}}^{(\text{DV})} \lesssim 0.06 \text{ fb}^{-1}$. Such cross sections are consistent with the constraints from displaced-jet searches quoted above. We have also confirmed, using the recasting tools associated with the MadAnalysis 5 [46] package, that each of these benchmarks is consistent with the LLP-search results [47] currently incorporated into the MadAnalysis database. We may therefore conclude that a significant number of both tumbler events and events involving DVs of any sort could potentially still be awaiting discovery at the LHC or at future colliders, even though no significant excess in such events has been observed to date. Although the above cross sections lie very close to the exclusion limits from displaced-jet searches, we also note that there are regions of our parameter space wherein $\sigma_{\text{eff}}^{(\text{DV})}$ lies even further below the bound from displaced-jet searches, tumblers still arise, and all additional constraints are satisfied.

Looking ahead, in order to assess what the results in Table IV portend in terms of the prospects for identifying a signal of new physics within the context of our model at the HL-LHC, we must take into account the relevant SM backgrounds. Fortunately, one of the advantages of searching for signal processes which lead to DVs is that these backgrounds are typically extremely low. One such background arises from SM processes which involve genuine DVs—for example, those associated with the decays of long-lived B - and K -mesons. However, the visible particles produced by these decays tend to be highly collimated whenever they are highly energetic as a result of the relatively small masses of the SM hadrons. By contrast, the particles produced by the decays of heavy LLPs into final states comprising significantly lighter particles are typically far less collimated. Indeed, this is the case for our example model when Δm_{21} and Δm_{10} are both $\mathcal{O}(100 \text{ GeV})$. Within this regime, cuts on variables which reflect the degree to which the visible particles produced at a DV are collimated—such as the uncertainty in the distance d_{BV} between the primary vertex and the DV [30,32] or, in the case of our example tumbler model, the angle between the three-momenta of the two

reconstructed jets—can be quite effective in reducing this background without a significant loss in the number of signal events.

After the contribution involving genuine DVs is suppressed in this way, the dominant contribution to the SM background in searches for displaced jets at the LHC is generally the one which arises as a consequence of multijet events in which poorly reconstructed tracks lead to the identification of spurious DVs [30,32]. Most such events arise from purely strong-interaction processes. Since both τ_1 and τ_2 satisfy $c\tau_n > \mathcal{O}(1 \text{ mm})$ for all of our benchmarks, the DVs that result from χ_1 and χ_2 decay are typically significantly farther than $0.1 \mu\text{m}$ away from the primary vertex. For displacements of this size, events involving additional primary vertices from pileup do not represent a significant background [30]. Thus, a rough estimate of the background event rate at the end of the full HL-LHC run can be obtained simply by scaling the expected number of background events obtained from searches using Run 2 data after the application of all relevant cuts by the ratio of the corresponding integrated luminosities.

The expected number of background events in any given displaced-jet search depends on the particular set of event-selection criteria employed, but the leading searches discussed above yield 0.1–0.7 background events [29,30] at an integrated luminosity of around 137 fb^{-1} . Thus, one would expect around 2.2–17.5 background events at the end of the full HL-LHC run. We also note that this background estimate is actually a conservative one, given that improvements in machine-learning approaches to LLP tagging have the potential to further reduce SM backgrounds without a significant loss in the signal-event rate [33]. By contrast, the signal efficiency obtained for the same cuts is typically around $\epsilon_S \sim 0.45\text{--}0.75$. Values within this range are obtained in both Ref. [29] and Ref. [30] for event topologies analogous to the one we consider here. Thus, given the results in Table IV, we see that a significant excess of DV-class events would be observed at the HL-LHC for all of our benchmarks. Moreover, for BM2 and BM3, this excess would include a substantial number of tumbler events. The observation of such an excess would clearly prompt significant additional investigation into how we might better probe the underlying physics responsible for this excess. It is toward this question that we now turn.

VI. DISTINGUISHING TUMBLERS VIA MASS RECONSTRUCTION

While we have shown that our model can give rise to a significant number of tumbler events at the LHC, we have also shown that it typically simultaneously gives rise to a far larger number of non-tumbler DV-class events—a substantial fraction of which likewise involve more than one DV. Indeed, any of the processes listed in Table III in which each decay chain involves only a single χ_1 or χ_2 particle gives rise to a pair of DVs. At this stage of the

analysis, such pairs of DVs are indistinguishable from tumblers. Thus, in this sense, our model not only gives rise to tumblers but also simultaneously gives rise to a “background” of non-tumbler events, each involving a pair of DVs which arise from decays within different chains. If a significant number of events involving multiple DVs is observed at the LHC either before or after the high-luminosity upgrade, it will therefore become imperative to develop methods of assessing whether or not a significant number of these events in fact involve tumblers.

This concern is not unique to our model alone. Indeed, there are also a variety of scenarios for physics beyond the SM in which events involving multiple DVs arise. These include SUSY models such as those constrained by the ATLAS and CMS searches in Refs. [38,39,41], as well as hidden-valley models [9] and other scenarios which give rise to emerging jets [5]. While tumblers can in fact arise within certain regimes in some of these models, many other models give rise to non-tumbler events exclusively. This then provides further motivation for developing methods of distinguishing between tumbler and non-tumbler events. Without doing so, one can not truly claim to have detected a tumbler signature.

Fortunately the distinctive kinematics associated with tumbler decay chains provides a basis on which we may discriminate between tumbler and non-tumbler events at colliders. In this section, we develop a set of event-selection criteria which are capable of efficiently discriminating between tumbler and non-tumbler events. In the process, we shall also investigate the extent to which the masses and lifetimes of the χ_n can be reconstructed from the kinematic and timing information provided by a collider detector.

A. Mass reconstruction

In order to distinguish between tumbler events and other events which involve multiple DVs, we employ an event-selection procedure which makes use of the distinctive kinematic structure associated with tumbler decay chains. This procedure follows from the observation that if two DVs in a given event arise from successive decays along the same decay chain, it is in principle possible to reconstruct the masses of the χ_n involved in that decay chain. That such an event-by-event mass reconstruction is possible for tumblers is itself noteworthy. Methods for reconstructing the masses of unstable particles in multi-step decay chains which terminate in invisible particles typically rely on the identification of features such as cusps [48–51], edges [51–65], or peaks [61,66,67] in the distributions of kinematic variables—features which emerge only in the aggregate, from a sizable population of events. By contrast, when the vertices in the decay chain are macroscopically displaced from each other and from V_p , as they are for a tumbler, additional information can be brought to bear in reconstructing the masses of the unstable particles.

The information we need in order to reconstruct the m_n for a tumbler includes the three-momenta of the four displaced jets produced by the decays of χ_1 and χ_2 , the three-momenta of the additional jets produced at the primary vertex, and the timing information supplied by the ECAL or timing layer concerning the time at which these jets exit the tracker. As discussed in Sec. II, these three-momenta, in conjunction with timing information, are sufficient to reconstruct the times t_P , t_S , and t_T and spatial locations \vec{x}_P , \vec{x}_S , and \vec{x}_T of the primary, secondary, and tertiary vertices. Taken together, these measurements are then sufficient to determine the velocities $\vec{\beta}_1 \equiv (\vec{x}_T - \vec{x}_S)/(t_T - t_S)$ and $\vec{\beta}_2 \equiv (\vec{x}_S - \vec{x}_P)/(t_S - t_P)$ of χ_1 and χ_2 , respectively.

Given these velocities, the m_n can then be determined in a straightforward manner. Approximating the quarks as massless and noting that the energy E_n and momentum \vec{p}_n of each χ_n are given by $E_n = \gamma_n m_n$ and $\vec{p}_n = \gamma_n m_n \vec{\beta}_n$, we find that the equations which represent four-momentum conservation at V_S may be written in the form

$$\begin{aligned} \gamma_2 m_2 &= \gamma_1 m_1 + |\vec{p}_q| + |\vec{p}_{\bar{q}}|, \\ \gamma_2 m_2 \vec{\beta}_2 &= \gamma_1 m_1 \vec{\beta}_1 + \vec{p}_q + \vec{p}_{\bar{q}}. \end{aligned} \quad (6.1)$$

Likewise, applying four-momentum conservation at V_T yields

$$\begin{aligned} \gamma_1 m_1 &= \gamma_0 m_0 + |\vec{p}_{q'}| + |\vec{p}_{\bar{q}'}|, \\ \gamma_1 m_1 \vec{\beta}_1 &= \vec{p}_0 + \vec{p}_{q'} + \vec{p}_{\bar{q}'}. \end{aligned} \quad (6.2)$$

Solving this system of equations for the three m_n , we obtain

$$\begin{aligned} m_2 &= \frac{|\vec{p}_q + \vec{p}_{\bar{q}} - \vec{\beta}_1(|\vec{p}_q| + |\vec{p}_{\bar{q}}|)|}{\gamma_2 |\vec{\beta}_1 - \vec{\beta}_2|}, \\ m_1 &= \frac{|\vec{p}_q + \vec{p}_{\bar{q}} - \vec{\beta}_2(|\vec{p}_q| + |\vec{p}_{\bar{q}}|)|}{\gamma_1 |\vec{\beta}_1 - \vec{\beta}_2|}, \\ m_0^2 &= m_1^2 - 2\gamma_1 m_1 [|\vec{p}_{q'}| + |\vec{p}_{\bar{q}'}| - \vec{\beta}_1 \cdot (\vec{p}_{q'} + \vec{p}_{\bar{q}'})] \\ &\quad + 2(|\vec{p}_{q'}||\vec{p}_{\bar{q}'}| - \vec{p}_{q'} \cdot \vec{p}_{\bar{q}'}). \end{aligned} \quad (6.3)$$

Were it possible to measure with arbitrary precision both the magnitude of the momentum of each jet in a tumbler event and the time at which each jet exits the tracker, it would be possible to reconstruct the m_n exactly from the relations in Eq. (6.3). In practice, of course, our ability to reconstruct these masses is limited by the precision with which the detector is capable of measuring these quantities. Nevertheless, provided that these uncertainties are sufficiently small, it is highly likely that the m_n values obtained when these reconstruction formulas are applied to the jets associated with a tumbler will satisfy certain basic self-consistency criteria. For example, these reconstructed m_n

values will be real, positive, and properly ordered in the sense that $m_2 > m_1 > m_0$.

By contrast, when the mass-reconstruction formulas in Eq. (6.3) are applied to the jets associated with a pair of DVs in the same event which do *not* arise from successive decays along the same decay chain, it is far less likely that they will yield a set of masses for the χ_n which satisfy these criteria. This consideration suggests that these mass-reconstruction formulas can be used in order to distinguish tumbler events from the far larger “background” of non-tumbler events involving multiple DVs which also arises in our model—and indeed arises generically in scenarios wherein the LLPs involved in the tumbler decay chain have identical quantum numbers.

In order to assess the extent to which we are able to distinguish tumbler events from other events involving multiple DVs in this way, we perform a Monte Carlo analysis. Our specific procedure is as follows. Using the MG5_aMC@NLO code package [45], and for each of our parameter-space benchmarks, we generate 100,000 events for the initial pair-production process $pp \rightarrow \phi^\dagger \phi$ at a center-of-mass energy $\sqrt{s} = 14$ TeV. This process overwhelmingly dominates the event rate for both tumbler-class and all relevant DV-class processes. The number of events in this sample is of course far larger than the expected event count for this pair-production process at the HL-LHC. Indeed, our goal at this stage of the analysis is simply to examine the detailed shapes of these distributions and thereby develop a nuanced understanding of how different event-selection criteria impact these shapes. It is therefore advantageous for us to consider a large population of events and a relatively narrow bin width for each m_n distribution. Once we have such an understanding, we shall return to assess the extent to which the m_n can be reconstructed with a population of events appropriate for near-future collider studies and a coarser set of bin widths.

After our events are generated, we then simulate the kinematics of the subsequent decay chains using our own Monte Carlo code. For each jet we record not only the magnitude and direction of its three-momentum vector, but also the time at which the jet exits the tracker. We work at the parton level and do not consider the effects of initial-state or final-state radiation, parton-showering, or hadronization. We determine the locations \vec{x}_S and \vec{x}_T of the secondary and tertiary vertices in each event from the momenta of the jets produced at these vertices using the parton-level vertexing algorithm described in Appendix. We likewise determine the location \vec{x}_P of the primary vertex from the momenta of the two jets produced by the prompt decays of ϕ and ϕ^\dagger at this vertex. Thus, while the beam spot at a collider like the HL-LHC has a characteristic spread of a few cm in the z direction and a time spread of around 200 ps, our procedure for reconstructing the primary vertex will effectively remove these uncertainties.

Of course, this parton-level vertexing procedure does not incorporate any of the uncertainties involved in a full track-based reconstruction of the locations of the primary or displaced vertices in the event. Moreover, it does not account for the measurement uncertainties in the momenta of the jets. Thus, in order to account for these uncertainties—which can be significant—when estimating the precision with which we might hope to measure the values of the m_n from tumbler data, we proceed as follows.

We account for the timing uncertainty by smearing the time at which each jet exits the tracker using a Gaussian smearing function with standard deviation σ_t . We likewise account for the uncertainty in the *magnitude* of the jet momenta by smearing the magnitude of each momentum vector according to a Gaussian smearing function whose standard deviation $\sigma_E(E_j)$ varies with the energy E_j of the jet. Since the jet-energy resolution of a collider detector also depends on the pseudorapidity η_j of the jet, we adopt a conservative approach and model our $\sigma_E(E_j)$ after the jet-energy resolution obtained in Ref. [68] for jets with $1.4 < \eta_j < 3.0$ in the endcap region rather than the barrel region of the CMS detector.

The uncertainties σ_η and σ_ϕ in the pseudorapidity and azimuthal angle that characterize the *direction* of each jet within a given event affect the reconstructed values of the m_n in two ways. The first is directly through \vec{p}_q , $\vec{p}_{\bar{q}}$, $\vec{p}_{q'}$, and $\vec{p}_{\bar{q}'}$ themselves in Eq. (6.3). The second is indirectly through their effect on the reconstructed vertex positions \vec{x}_p , \vec{x}_s , and \vec{x}_t , which in turn affects the reconstructed LLP velocities $\vec{\beta}_1$ and $\vec{\beta}_2$. Since the CMS detector is capable of measuring the directions of the momentum vectors of hadronic jets with excellent precision [68], the first effect turns out to be subleading in terms of its effect on the m_n in comparison with the effect of jet-energy smearing. By contrast, the second effect can have a more significant impact on the m_n . Indeed, σ_η and σ_ϕ can dominate the uncertainty in $\vec{\beta}_1$ and $\vec{\beta}_2$ when σ_t is small.

Our method for simulating the effect of these uncertainties shall be the following. Since σ_E dominates the uncertainty in the m_n that arises directly from the jet momenta, we shall simply take $\sigma_\eta = \sigma_\phi = 0$ in what follows. However, in order to account for the effect of these angular uncertainties and other uncertainties which enter into the track-based reconstruction of DVs at a real collider detector, we also shift each of the three vertex positions \vec{x}_p , \vec{x}_s , and \vec{x}_t that we obtain from our fitting procedure by an independent random offset vector. The magnitude of this offset vector is distributed according to a single-sided Gaussian function with standard deviation σ_r , while its direction is distributed spherically uniformly. Since the estimated uncertainty in the vertex displacements for the CMS detector after the HL-LHC upgrade is roughly $\mathcal{O}(10\text{--}30\ \mu\text{m})$ [7], we take $\sigma_r = 30\ \mu\text{m}$ in what follows.

TABLE V. The fraction of the events with the topology given in the fourth line of Table III (i.e., with decay chains $\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$ and $\phi \rightarrow \chi_0$) in which the last unstable particle decays within each layer of the detector for each of our benchmarks BM1–BM4.

Benchmark	Tracker	ECAL+HCAL	Muon chamber	Outside detector
BM1	0.56	0.26	0.15	0.03
BM2	0.74	0.19	0.06	0.01
BM3	0.77	0.17	0.06	0.01
BM4	1.00	0.00	0.00	0.00

In order to extract a set of values for the m_n from a given sample of events, as well as an estimate of the uncertainties in these values, we proceed as follows. We begin by requiring that the decays of all unstable dark-sector particles in the event occur within the tracker region of our hypothetical detector. Modeling this detector after the CMS detector, we take this region to be a cylinder of radius $r = 1.161\ \text{m}$, centered at the interaction point $z = 0$ and extending longitudinally within the range $-2.5\ \text{m} < z < 2.5\ \text{m}$, whose axis of symmetry runs along the beam. We note that events which satisfy this requirement necessarily involve a significant number of energetic jets—including two highly energetic prompt jets from the decays of ϕ and ϕ^\dagger —and typically also significant \cancel{E}_T . The overwhelming majority of such events therefore satisfy one or more of several Level-1 triggers appropriate for a detector in high-luminosity collider environment [69].

In order to assess the impact of this requirement on our results, we focus on the events which have the event topology given in the fourth line of Table III—i.e., events wherein the two decay chains are $\phi \rightarrow \chi_2 \rightarrow \chi_1 \rightarrow \chi_0$ and $\phi \rightarrow \chi_0$. Events of this sort, which involve a single tumbler but no additional unstable particles on the other side of the event, provide a the clearest picture of where these decays tend to occur. In Table V, for each of our parameter-space benchmarks, we provide the fraction of events in our Monte Carlo sample with this event topology in which the χ_1 decays within each layer of the detector. We observe that while a non-negligible fraction of these particles decay outside the tracker for all of these benchmarks except BM4, which has a far smaller value of τ_1 than the other three benchmarks, the χ_1 particle decays within the tracker the majority of the time. By contrast, τ_2 is sufficiently short for all of our benchmarks that the probability for χ_2 to decay outside the tracker is negligible. The results shown in Table V indicate that the requirement that all unstable LLPs in the event decay within the tracker, rather than elsewhere within the detector, will not have a significant impact on our results. Moreover, they also indicate that the fraction of events in which χ_1 escapes the detector entirely before decaying is quite small for all of our benchmarks.

We also require that the event contain at least two DVs. We compute the time t_i at which each such vertex V_i occurred from the momentum and timing information obtained for the pair of displaced jets produced at that vertex. For each combination of DVs V_i and V_j in the event which are appropriately time-ordered, in the sense that $t_i < t_j$, we reconstruct a set of m_n values using Eq. (6.3). We then check whether this set of m_n values, taken together with the corresponding values of $|\vec{\beta}_1|$, $|\vec{\beta}_2|$, and the magnitude of the three-momentum vector \vec{p}_0 obtained from Eq. (6.2), satisfy the following criteria, to which we shall henceforth refer as our *reconstruction criteria*:

- m_1 and m_2 are real and positive;
- m_0^2 is real;
- $|\vec{p}_0|$ is real and positive;
- $0 < |\vec{\beta}_n| < 1$ for $n = 1, 2$;
- $m_2^2 > m_1^2 > m_0^2$.

For reasons to be discussed shortly, we shall not require that $m_0^2 > 0$ at this stage of the analysis. If any appropriately time-ordered combination of DVs in the event yields a set of masses which satisfy these criteria, we retain the event; if not, we reject it. If multiple combinations of DVs within the same event satisfy all of these criteria, we take the set of m_n for the combination which yields the largest value of m_2 to be the set of m_n for the event.

In order to illustrate the effect of these cuts, we shall begin by focusing on the reconstruction of m_1 . In Fig. 6, we show the distribution of reconstructed m_1 values for the set of events which survive these cuts for each of our parameter-space benchmarks. The histogram in each panel of the figure is obtained by binning these m_1 values into bins of width $\Delta m_n = 5$ GeV. The blue portion of each histogram bar represents the contribution to that bin from non-tumbler processes, whereas the orange portion represents the contribution from processes which involve tumblers. From top to bottom, the rows in the figure correspond to our parameter-space benchmarks BM1–BM4. The dashed black vertical line in each panel indicates the actual value of m_1 for the corresponding benchmark. The results shown in the left, center, and right columns correspond respectively to the values $\sigma_t = 30$ ps, $\sigma_t = 5$ ps, and $\sigma_t = 0.01$ ps for the timing uncertainty of the detector. The first of these σ_t values represents the timing uncertainty associated with the barrel timing layer to be installed within the CMS detector as part of the HL-LHC upgrade. The second is a value chose to reflect a moderate improvement in this timing uncertainty, while the third is an extremely small value representative of the regime in which jet-energy and vertex-position smearing dominates the uncertainty in the mass reconstruction. Since the efficiency of the cuts depends on the benchmark and varies with σ_t , the scale of the vertical axis has been varied from panel to panel in order to facilitate comparison between the distributions.

First, we observe from Fig. 6 that the number of residual non-tumbler events is still quite significant even after the

imposition of these preliminary cuts. Moreover, we observe that this distribution has a well-defined shape that peaks at low values of m_1 and falls off rapidly as m_1 increases. By contrast, the m_1 distribution for the tumbler events exhibits a well-defined peak centered around the actual value of m_1 , as well as an additional population of events with m_1 values well below this peak. This additional population of events arises in part due to smearing effects and in part due to the combinatorial background which arises from incorrect identifications of the vertices V_S and V_T in events which contain more than two DVs. The relative size of the peak in the m_1 distribution for the tumbler events at low σ_t is primarily controlled by $P_{\phi 210}$. Indeed, we observe that this peak is more pronounced for BM3, which has by far the largest value of $P_{\phi 210}$, than for our other three benchmarks.

The presence of this peak in the m_1 distribution is a unique and distinctive feature of tumbler events. As we shall see, similar peaks appear in the distributions of m_0 and m_2 for tumbler events as well. An observation of these peaks, taken together, would constitute compelling evidence for tumblers. It is in this way, then, that our mass-reconstruction procedure furnishes a method through which tumblers can unambiguously be detected.

We also observe from Fig. 6 that as σ_t decreases, the peak in the tumbler distribution becomes both narrower and more pronounced for all of our benchmarks. Indeed, this is to be expected, since increasing σ_t renders the reconstructed values of t_p , t_S , and t_T less reliable. However, a greater reduction in timing uncertainty is required to resolve this peak for some of our benchmarks than for others. For example, the peak obtained for BM4 remains effectively washed out even for $\sigma_t = 5$ ps. We can make sense of these differences in sensitivity to σ_t by comparing the lifetimes τ_1 and τ_2 quoted for each of our benchmarks in Table II to the value of σ_t itself. For BM1 and BM2, $\tau_1 \gg \tau_2 \sim \mathcal{O}(100$ ps), and thus the effect of the timing uncertainty on the times t_S and t_T reconstructed for the DVs in a tumbler event will be negligible for either of these benchmarks when $\sigma_t \ll 100$ ps. By contrast, for BM3 and BM4, $\tau_2 \sim \mathcal{O}(10$ ps), which implies that the effect of the timing uncertainty on t_S will only be negligible when $\sigma_t \ll 10$ ps. Furthermore, for BM4, $\tau_1 \sim \mathcal{O}(100$ ps) is also quite small, and thus the timing uncertainty has a non-negligible impact on t_T as well unless $\sigma_t \ll 100$ ps. As a result, the reconstructed value of m_1 is more sensitive to the value of σ_t for BM4 than they are for BM3, and are more sensitive to this value for BM3 than they are for BM1 and BM2.

In order to further suppress the contribution from non-tumbler events, we shall impose one additional cut on the data. In particular, in addition to the criteria described above, we shall also impose one additional reconstruction criterion:




FIG. 6. The distribution of values of the mass m_1 for the sample of Monte Carlo events described in the text, as reconstructed from tumbler kinematics. The orange portion of each histogram bar represents the contribution from tumbler events, while the blue portion represents the contribution from events with multiple DVs which do not involve a tumbler. From top to bottom, the rows in the figure correspond to the parameter-space benchmarks BM1–BM4 defined in Table I. The dashed black vertical line in each panel indicates the actual value of m_1 for the corresponding benchmark. The results shown in the left, center, and right columns correspond respectively to the values $\sigma_t = 30$ ps, $\sigma_t = 5$ ps, and $\sigma_t = 0.01$ ps for the timing uncertainty of the detector. Since the efficiency of the cuts depends on the benchmark and varies with σ_t , the scale of the vertical axis has been varied from panel to panel in order to facilitate comparison between the distributions.

- $m_0^2 > 0$.

We have separated out this particular criterion from the others because it merits special attention. In particular, as we shall demonstrate, not only does requiring that $m_0^2 > 0$ induce a dramatic enhancement in the ratio of tumbler to non-tumbler events, but it also gives rise to an additional feature in the distribution of reconstructed m_1 values—a feature which reveals additional information about the mass spectrum of the χ_n , and in particular about the mass splitting Δm_{10} .

In order to quantify the effect of the $m_0^2 > 0$ criterion on the ratio of tumbler to non-tumbler events for each of our four parameter-space benchmarks, in Fig. 7 we plot the ratio of the number N_T of tumbler events to the number N_{NT} of non-tumbler events obtained for each of our parameter-space benchmarks after cuts as a function of σ_t . The dash-dotted curves represent the N_T/N_{NT} ratios obtained after the imposition of all of our event-selection criteria *except* the $m_0^2 > 0$ criterion. By contrast, the solid curves represent the




FIG. 7. The ratio of the number N_T of tumbler events to the number N_{NT} of non-tumbler events for each of our parameter-space benchmarks, shown as a function of the timing uncertainty σ_t . The dash-dotted curves in each panel represent the corresponding efficiencies obtained *without* imposing the $m_0^2 > 0$ criterion, whereas the solid curves represent the corresponding efficiencies obtained with the $m_0^2 > 0$ criterion included. The vertical arrows in each case therefore indicate the improvements induced by imposing the $m_0^2 > 0$ cut.

N_T/N_{NT} ratios obtained after the $m_0^2 > 0$ criterion is also imposed.

It is evident from Fig. 7 that the imposition of the $m_0^2 > 0$ criterion has a significant impact on N_T/N_{NT} . When σ_t is relatively large, as on the right side of this figure, this enhancement factor is already significant for our first three benchmarks, even up to $\sigma_t = 30$ pb. By contrast, as σ_t decreases (towards the left side of this figure), this ratio is enhanced even further, ultimately reaching a factor of ~ 10 for all of the benchmarks. The only exception to this behavior arises for BM4. For BM4, the value of σ_t has a proportionally greater effect on the times reconstructed for the DVs and the difference between τ_1 and τ_2 is far smaller than for our other benchmarks. As a result of these differences, the effect of smearing σ_t is more likely to result in a set of reconstructed masses which fail our reconstruction criteria for BM4 than it is for our other three benchmarks.

We now turn to discuss the impact of the $m_0^2 > 0$ criterion on the *shapes* of the m_n distributions obtained from our mass-reconstruction procedure, and in particular on the shape of the m_1 distribution—the distribution on which this criterion has the greatest impact. Indeed, since the presence of identifiable peaks in each of the three reconstructed m_n distributions is the characteristic feature that distinguishes a population of tumbler events from a

population of non-tumbler events, the shapes of these distributions are of crucial importance.

The shapes of these distributions also allow us to determine the masses of the LLPs involved in the tumbler. In order to assess the precision with which this can be done, we need a method of estimating the width of the peak in the corresponding mass distribution. We shall do this in the following way. We begin by constructing a template for the non-tumbler contribution to each m_n distribution after the application of our event-selection criteria. We construct each such template by performing a smoothing procedure on the non-tumbler contribution to the m_n distribution obtained from an additional sample of Monte Carlo events—a smoothing procedure wherein we replace the number of events in each histogram bin with the mean value of the event counts in all bins whose central m_n values are within 25 GeV of the central m_n value for that bin. We then subtract this template from the corresponding m_n distribution in order to obtain the contribution to the m_n distribution from the tumbler events alone. We then perform a fit of this “background-subtracted” m_n distribution to the rescaled Gaussian function

$$f(m_n) = \frac{N_{m_n}}{\sqrt{2\pi\sigma_{m_n}^2}} \exp\left[-\frac{(m_n - \langle m_n \rangle)^2}{2\sigma_{m_n}^2}\right]. \quad (6.4)$$

We take the values of $\langle m_n \rangle$ and σ_{m_n} as our best estimates for m_n and its uncertainty. While more sophisticated modeling of the shape of the mass peak would of course improve upon these results, this procedure provides a reasonably reliable indicator of the extent to which one might hope to extract a meaningful measurement of each m_n for a given set of model parameters at the LHC or at future colliders.

In Fig. 8, we display the m_1 distributions obtained for our benchmarks after the application of all of our event-selection criteria, *including* the $m_0^2 > 0$ criterion. Thus, all differences between the m_1 distribution shown in each panel of this figure and the distribution shown in the corresponding panel of Fig. 6 are solely due to the effect of this criterion. The $\langle m_1 \rangle$ and σ_{m_1} values we obtain from our fitting procedure for the distribution shown in each panel are also indicated.

We observe that the non-tumbler contribution to each of the m_1 distributions shown in Fig. 8 is significantly reduced relative to the corresponding distribution in Fig. 6. However, somewhat surprisingly, we also see that each of these distributions now manifests a visible dip or trough at a particular reconstructed value of m_1 well below this peak. The origin of this dip can be understood as follows. First, we see from Eq. (6.3) that events which fail to satisfy the $m_0^2 > 0$ criterion are events for which

$$m_1^2 - 2m_1 E_{jj}^* + m_{jj}^2 \leq 0, \quad (6.5)$$

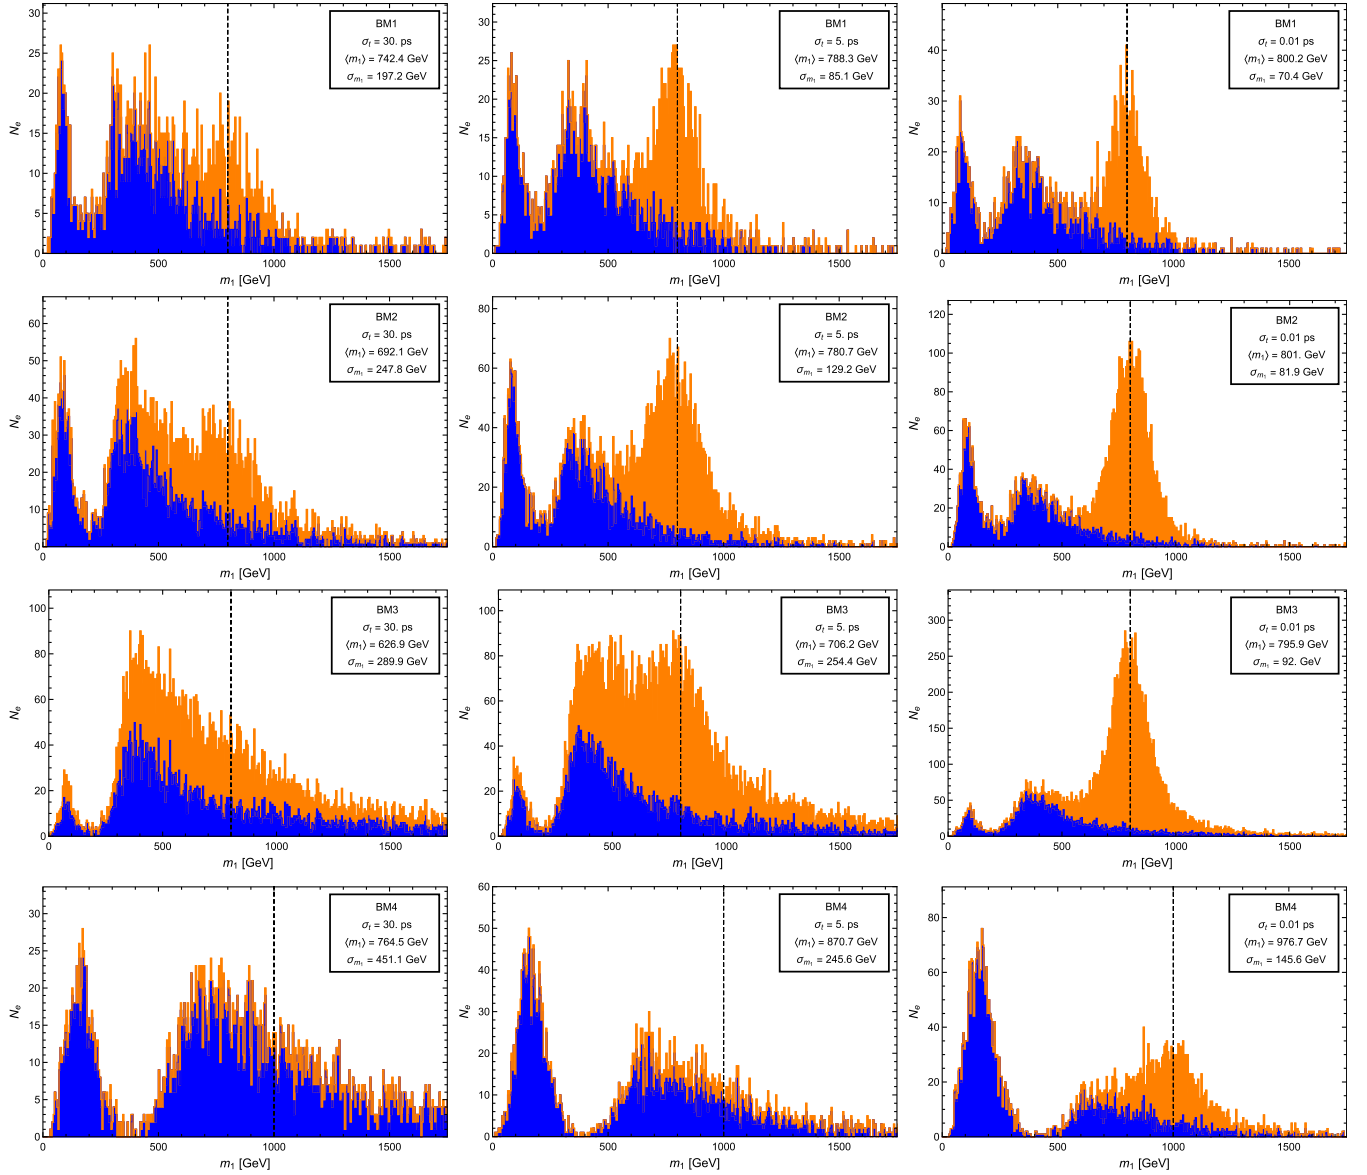


FIG. 8. Same as Fig. 6, but after the imposition of the additional $m_0^2 > 0$ reconstruction criterion.

where we have used the fact that the center-of-mass energy reconstructed for the $q\bar{q}'$ system is given by $E_{jj}^* \equiv \gamma_1[|\vec{p}_{q'}| + |\vec{p}_{\bar{q}'}| - \vec{\beta}_1 \cdot (\vec{p}_{q'} + \vec{p}_{\bar{q}'})]$ and the fact that the invariant mass of this system is given by $m_{jj}^2 = 2(|\vec{p}_{q'}||\vec{p}_{\bar{q}'}| - \vec{p}_{q'} \cdot \vec{p}_{\bar{q}'})$ in order to write this condition more compactly. Thus, for any particular values of E_{jj}^* and m_{jj} , the range of reconstructed m_1 values excluded by the $m_0^2 > 0$ criterion is

$$E_{jj}^* - \sqrt{(E_{jj}^*)^2 - m_{jj}^2} \leq m_1 \leq E_{jj}^* + \sqrt{(E_{jj}^*)^2 - m_{jj}^2}. \quad (6.6)$$

We note that this range of excluded m_1 values always contains the point $m_1 = E_{jj}^*$.

We also observe that constraints which follow from standard three-body-decay kinematics restrict the true values of E_{jj}^* and m_{jj} to lie within the respective ranges $0 \leq m_{jj} \leq m_1 - m_0$ and $(m_1^2 - m_0^2)/(2m_1) \leq E \leq m_1 - m_0$. Of course, the reconstructed values of E_{jj}^* and m_{jj} will in general differ from these true values due to timing, jet-energy, and vertex-position smearing, and can in principle lie outside these ranges. However, in the regime in which σ_t is negligible compared to τ_1 and τ_2 , we find that the vast majority of reconstructed values for E_{jj}^* and m_{jj} lie within or only slightly outside these ranges. For all of our parameter-space benchmarks, we note that the range of kinematically allowed E_{jj}^* values is fairly narrow. For BM1–BM3, this range is $175 \text{ GeV} \leq E_{jj}^* \leq 200 \text{ GeV}$; for BM4, this range is $320 \text{ GeV} \leq E_{jj}^* \leq 400 \text{ GeV}$. As a

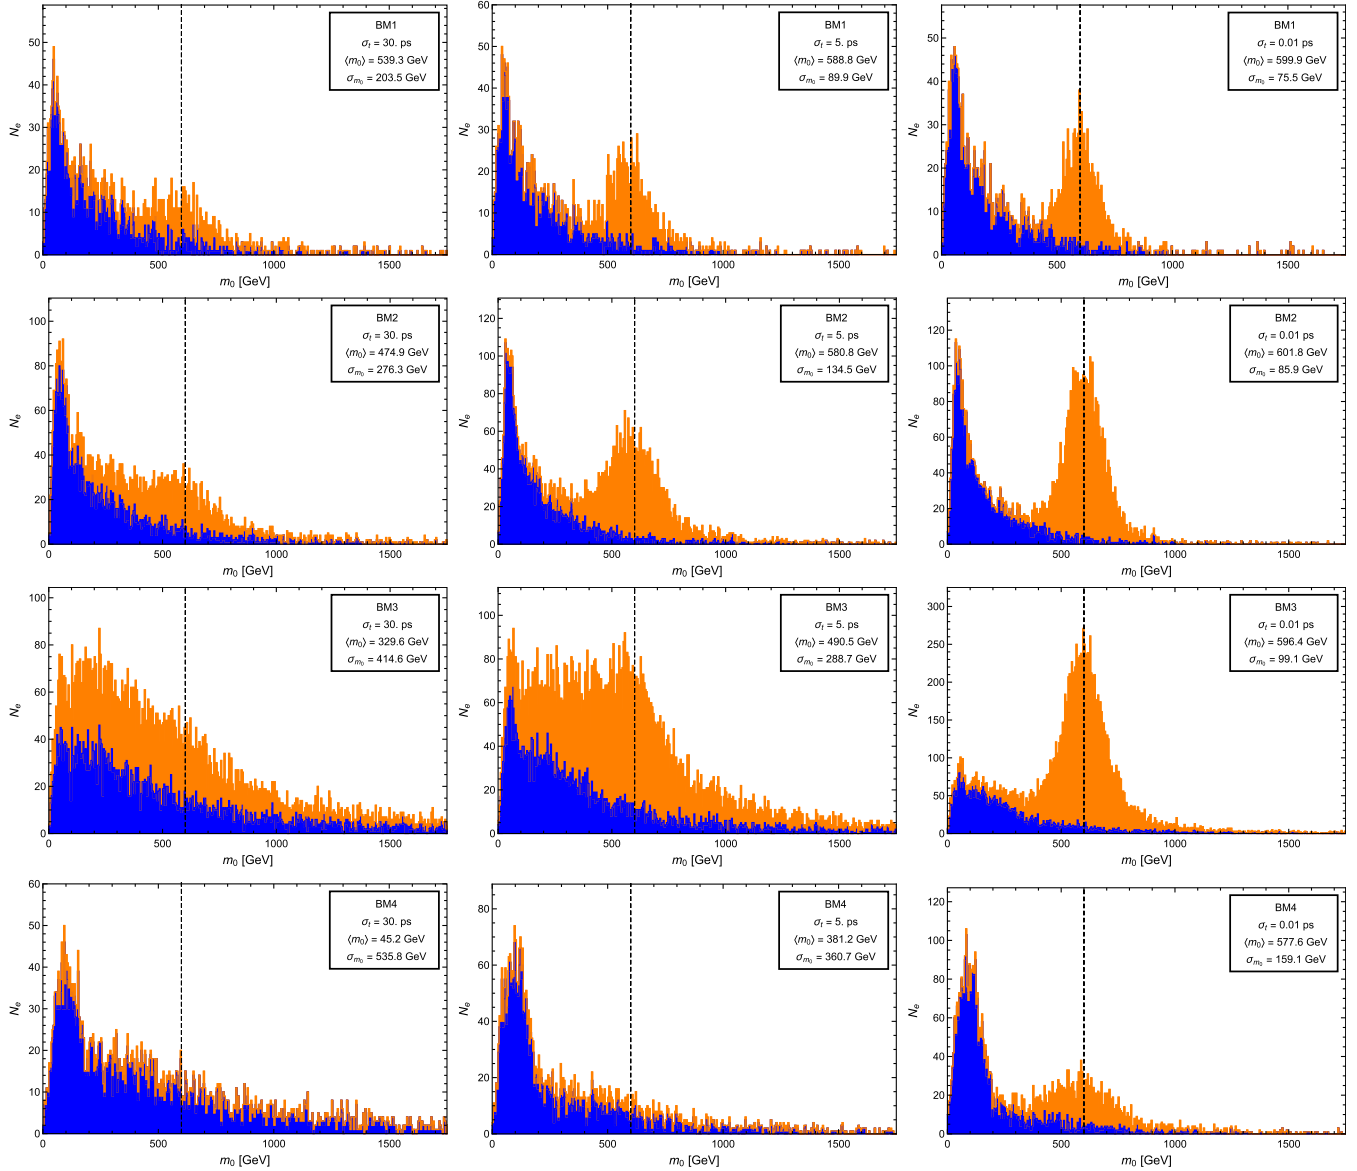


FIG. 9. Same as Fig. 8, except that the distributions shown are for m_0 rather than m_1 .

result, when the reconstructed value of m_1 for an event lies within this narrow range of E_{jj}^* values, Eq. (6.6) implies that the event will typically be excluded. Indeed, we observe a dramatic suppression in each of the distributions shown in Fig. 8 across the corresponding range of m_1 values.

It is worth remarking that this dip in the m_1 distribution arises solely as a consequence of the decay kinematics at the final vertex V_T along the tumbler decay chain. Thus, the kinematic considerations which lead to the dip are insensitive to the full structure of that decay chain. We would therefore expect the contribution to the m_1 distribution from non-tumbler events in which a χ_1 particle appears in either one or both of the decay chains to exhibit a similar dip. Indeed, we observe that a dip appears in both tumbler and non-tumbler contributions to the m_1 distributions in

Fig. 8. It is also worth remarking that the location and width of the dip provide additional information about the mass spectrum of the χ_n . Indeed, we have seen that both E_{jj}^* and m_{jj} are bounded from above by Δm_{10} . Thus, in principle, correlations between the properties of the dip and the locations of the tumbler peaks in the m_n distributions can be exploited to improve the precision with which the m_n can be measured.

We also observe from Fig. 8 that the width σ_{m_1} of the tumbler peak obtained for each benchmark depends quite sensitively on σ_t . When σ_t is fairly large, as shown in the left and center columns of this figure, timing uncertainty tends to dominate the widths of the peaks in the tumbler distributions. By contrast, when σ_t is sufficiently small, as in the right column of this figure, the widths of these peaks are instead dominated by σ_E and σ_r . The value of σ_t at

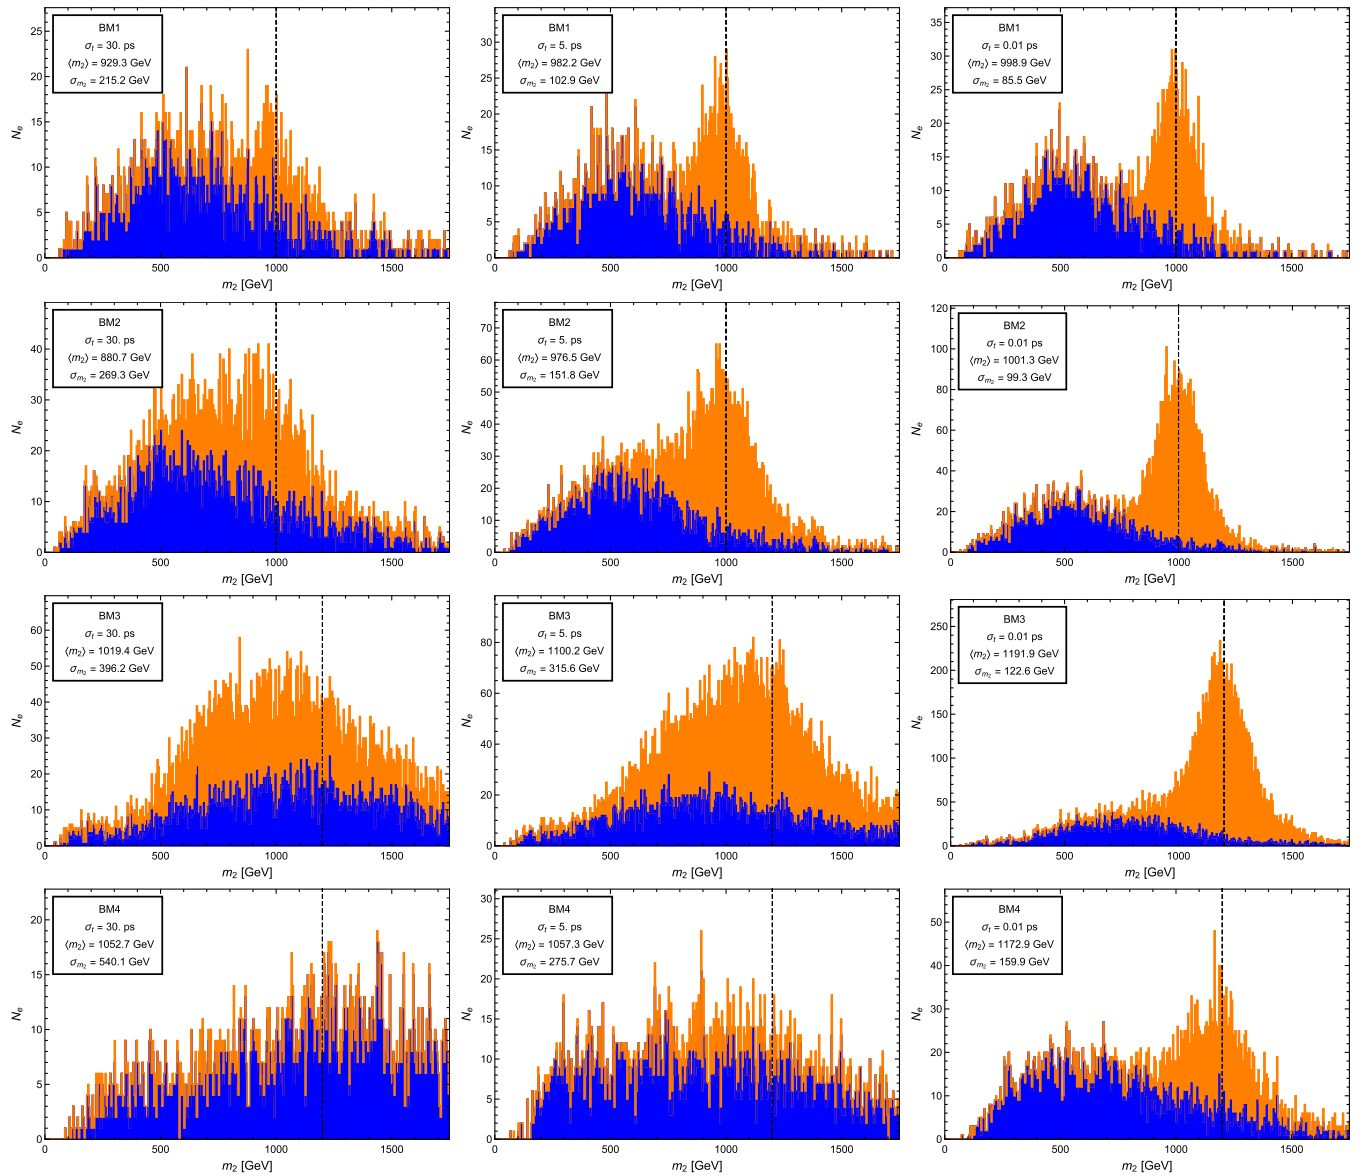


FIG. 10. Same as Fig. 8, except that the distributions shown are for m_2 rather than m_1 .

which this transition occurs for each of our benchmarks depends once again on τ_1 and τ_2 . Nevertheless, it is clear that the identification of a tumbler peak in the m_1 distribution will be extremely challenging with a timing resolution on the order of the $\sigma_\tau = 30$ ps that the CMS barrel timing layer will be able to provide at the beginning of the upcoming HL-LHC run. However, it is also clear that a reduction in timing uncertainty by even a factor of a few relative to this value would significantly enhance the capabilities of the HL-LHC or future colliders—both in terms of distinguishing tumblers from other signatures of new physics involving multiple DVs and in terms of extracting information about the mass spectrum of the particles involved. Indeed, the results shown in Fig. 8 are an indication that we are on the doorstep of being able to probe the underlying physics which gives rise to DVs at a much deeper level.

Thus far, we have focused on the reconstruction of the mass m_1 . In Figs. 9 and 10, we show the corresponding distributions of reconstructed m_0 and m_2 values for our four benchmarks, respectively, after the application of our event-selection criteria, including the $m_0^2 > 0$ criterion. As with the m_1 distributions, there are no significant discernible peaks when σ_τ is larger than $\mathcal{O}(1-5)$ ps. This is true for all benchmarks. However, as σ_τ decreases, a discernible peak begins to appear in both the m_0 and m_2 distributions, ultimately becoming higher and narrower as σ_τ drops. Moreover, for each benchmark, these peaks are centered around the true values of the corresponding masses. However, unlike the distributions shown in Fig. 8, the distributions in Figs. 9 and 10 do not exhibit a discernible dip at any particular value of the corresponding reconstructed m_n .

Taken together, the results shown in Figs. 8–10 attest that our mass-reconstruction procedure is quite effective in discriminating between tumbler and non-tumbler events, provided that the timing uncertainty is sufficiently small that the peaks in the m_n distributions can be resolved. On the one hand, it is clear from these figures that conclusively identifying tumblers at the HL-LHC with the $\sigma_t \approx 30$ ps timing resolution the CMS timing layer is anticipated to provide would prove challenging indeed. On the other hand, it is also clear that a moderate reduction in timing uncertainty from $\sigma_t \approx 30$ ps to $\sigma_t \approx 5$ ps would have a dramatic effect on our ability to probe the underlying structure of the decay chains that give rise to events involving multiple DVs. As we demonstrated in Sec. V, a robust excess in the relevant detection channels could yet be observed at the LHC. If such an excess is in fact observed, improvements in timing precision, in conjunction with event-selection procedures like the one we have developed here, will play a pivotal role in determining whether or not this excess arises as a consequence of successive decays within the same decay chain.

B. Lifetime reconstruction

We now assess the degree to which we can likewise measure the respective lifetimes τ_1 and τ_2 of the unstable LLPs involved in the tumbler decay chain. For any given tumbler, the proper-time intervals t_1 and t_2 between the production and decay of each of these particles are given by $t_1 = (t_T - t_S)/\gamma_1$ and $t_2 = (t_S - t_P)/\gamma_2$, where $\gamma_n \equiv (1 - |\vec{\beta}_n|)^{-1/2}$ is the usual relativistic factor. In order to estimate the proper lifetime τ_n of each particle from a given sample of events, we first select events which satisfy the same criteria we imposed in our mass-reconstruction analysis. We then define $N_n(t)$ to represent the number of events in the sample for which $t_n > t$. We then perform a least-squares fit of the function $f(t) = N_n(0) \exp(-t/\tau_n)$ to the events in the sample and interpret the value of τ_n as our estimate for the proper lifetime of χ_n . Since the goodness-of-fit statistic for this nonlinear fit is more sensitive to deviations in which t is small and $N_n(t)$ is large, the resulting value of τ_n is typically insensitive to the small, residual contribution to $N_n(t)$ at large t from non-tumbler events which nevertheless survive our mass-reconstruction cuts.

In Fig. 11, we show the results of such a fit for the parameter-space benchmarks defined in Table I. The orange and blue histograms in each panel respectively represent the $N_1(t)$ and $N_2(t)$ distributions obtained for a Monte Carlo data sample that once again initially consists of 100,000 events prior the imposition of our event-selection criteria. However, only the $N_1(t)$ and $N_2(t)$ for events which pass all of these cuts are included in the histograms. The thick orange and blue curves represent the exponential-decay functions obtained for our best-fit values of $c\tau_1$ and $c\tau_2$, respectively. From top to bottom, the rows

in the figure correspond to our parameter-space benchmarks BM1–BM4. The results shown in the left, center, and right columns of Fig. 11 once again correspond respectively to the timing uncertainties $\sigma_t = 30$ ps, $\sigma_t = 5$ ps, and $\sigma_t = 0.01$ ps.

We begin by noting that each $N_n(t)$ distribution shown in Fig. 11 clearly includes contributions from two distinct populations of events. The first of these populations, which is far larger than the second and dominates $N_n(t)$ when ct is small, comprises genuine tumbler events. The second population, which includes events with much higher ct values, consists primarily of residual non-tumbler events. Since this second population is quite small, it does not have a dramatic impact on the best-fit value of the corresponding $c\tau_n$.

The results shown in Fig. 11 demonstrate that for relatively large σ_t values, the accuracy with which these lifetimes can be measured differs among the different benchmarks. For example, the extent to which our fitting procedure overestimates the value of τ_2 for BM3 and BM4 is significant, whereas this effect is less severe for BM1 and BM2. This is once again primarily a reflection of the fact that τ_2 is far shorter for BM3 and BM4 than it is for these other benchmarks, and hence the effect of timing uncertainty on the results for BM3 and BM4 becomes significant at a far lower value of σ_t .

Somewhat counterintuitively, however, we also observe that our fit systematically underestimates the value of τ_1 for BM1–BM3 by as much as a factor of two when σ_t is small. This is a consequence of τ_1 being sufficiently large for these benchmarks that a small but non-negligible fraction of the χ_1 particles produced by χ_2 decays themselves decay outside the timing layer. Since events in which these χ_1 particles decay outside the timing layer are of course not included in any of our event samples, the $N_1(t)$ distribution is slightly skewed toward lower lifetimes. Thus, as σ_t decreases, the best-fit value of $c\tau_1$ approaches a value slightly below the actual proper decay length. This effect is not particularly significant for BM4, however, since τ_1 is far shorter and the fraction of events in which χ_1 escapes the detector before it decays is therefore far smaller. That said, we emphasize that reasonably reliable measurements of both τ_1 and τ_2 can nevertheless be made for all four of our benchmarks when $\sigma_t = 0.01$ ps, even for the simple, physically motivated functional fit we have performed here. An alternative functional fit which accounts for this finite-volume effect could yield even better estimates of the LLP lifetimes.

C. Tumbler searches with limited event counts

We now consider the extent to which we are able to resolve the characteristic tumbler peaks in the distributions of the reconstructed m_n values given a far smaller number of events—a number which might realistically be obtained at the HL-LHC or at other, near-future colliders. For this

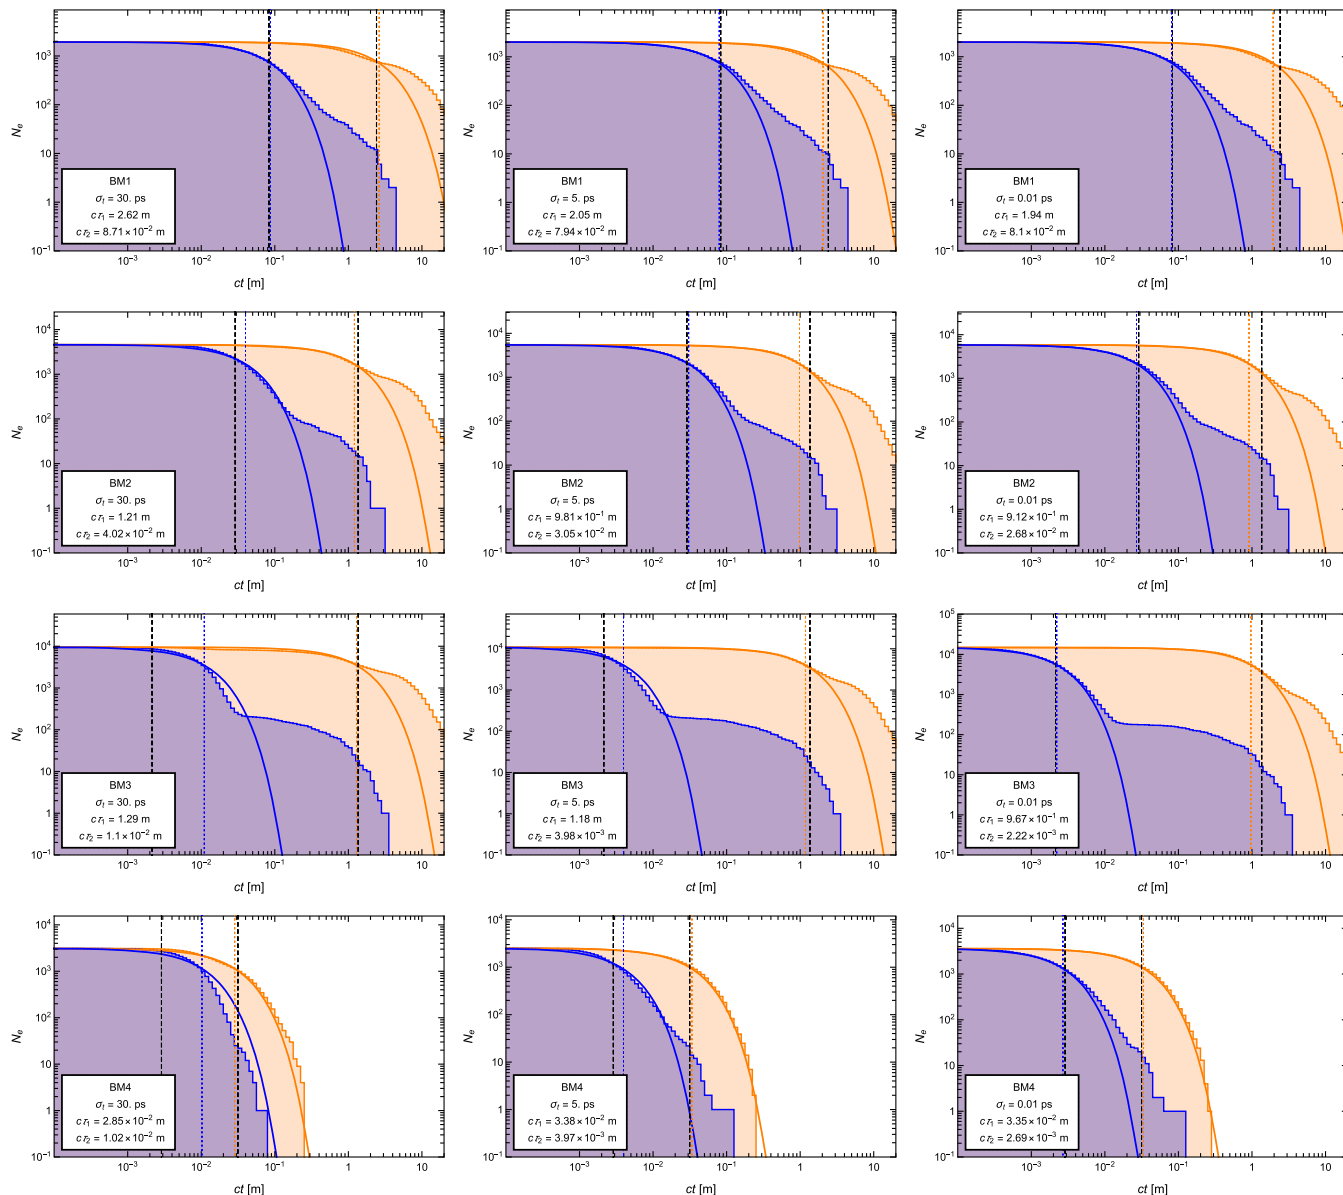


FIG. 11. Distributions of the number of events $N_1(t)$ (orange histogram) and $N_2(t)$ (blue histogram) for which the corresponding LLP χ_1 or χ_2 has not yet decayed a proper time t after it was initially produced, displayed as a function of the corresponding proper decay distance ct . From top to bottom, the rows in the figure correspond to the parameter-space benchmarks BM1–BM4 defined in Table I. The results shown in the left, center, and right columns correspond respectively to the values $\sigma_t = 30$ ps, $\sigma_t = 5$ ps, and $\sigma_t = 0.01$ ps for the timing uncertainty of the detector. Exponential-decay curves constructed using the best-fit values of ct_1 (thick orange curve) and ct_2 (thick blue curve) are also shown in each panel. The dotted orange and blue vertical lines correspond to the best-fit values of ct_1 and ct_2 , respectively, while the dashed black vertical lines indicate the actual values of these proper decay lengths. The best-fit values of ct_1 and ct_2 are also quoted in the box in the lower left corner of each panel.

purpose, we shall consider a hypothetical collider essentially identical to the HL-LHC. The two general-purpose detectors at this collider are each assumed to be equipped with a barrel timing layer with timing uncertainty σ_t , but to be otherwise similar in design and performance to the CMS detector. We assume an integrated luminosity $\mathcal{L}_{\text{int}} = 6000 \text{ fb}^{-1}$ in each detector—an integrated luminosity equal to twice that anticipated for the HL-LHC over its full run.

Thus, the total event count for tumbler events before cuts is taken to be $2\mathcal{L}_{\text{int}}\sigma^{(T)}$, and the total number of non-tumbler events including at least one DV is calculated in an analogous manner.

In Fig. 12, we show the distribution of reconstructed m_1 values for a Monte Carlo data set consisting of the expected number of events for each of our parameter-space benchmarks at such a pair of collider detectors.

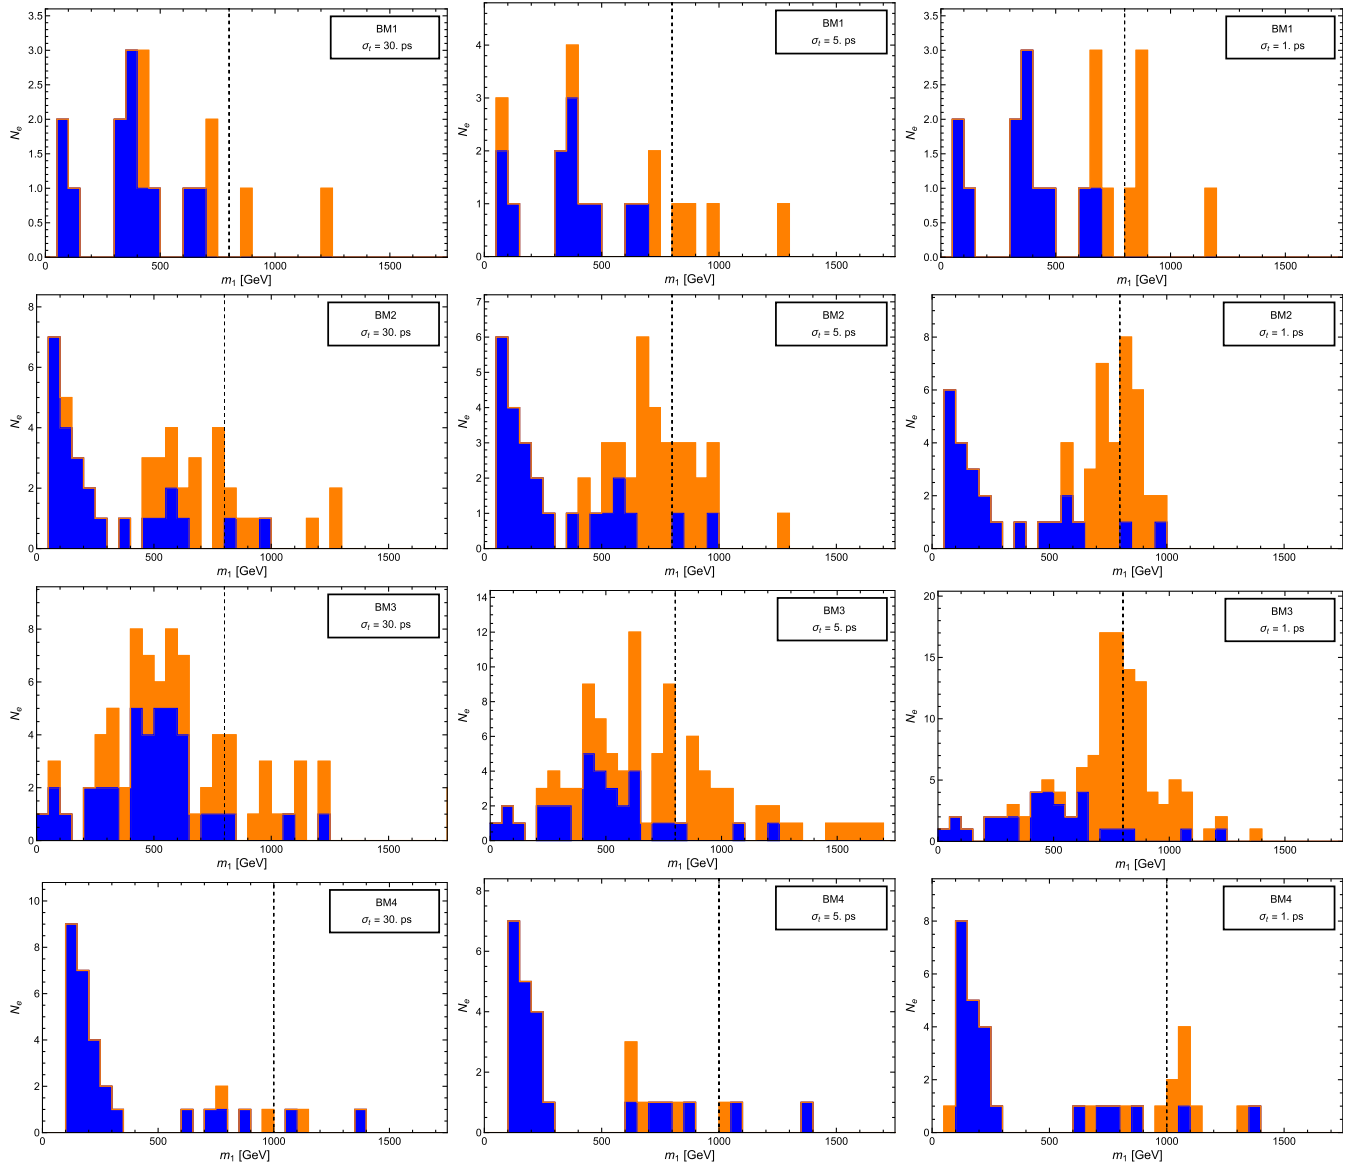


FIG. 12. Same as Fig. 8, but for a smaller sample of Monte Carlo events. In particular, the numbers of both tumbler and non-tumbler DV events included in each data sample before any cuts are applied are equal to the expected numbers of events at a collider essentially identical to the HL-LHC, but with an integrated luminosity twice that anticipated for the full HL-LHC run. The results shown in the left, center, and right columns correspond respectively to the values $\sigma_t = 30$ ps, $\sigma_t = 5$ ps, and $\sigma_t = 1$ ps for the timing uncertainties of the two general-purpose detectors at this collider, and we have adopted a bin width of 50 GeV when constructing each histogram.

Only events which survive all of our cuts—including the $m_0 > 0$ criterion—are included in each distribution shown. From top to bottom, the rows in the figure correspond to the parameter-space benchmarks defined in Table I. The results shown in the left, center, and right columns correspond respectively to the timing uncertainties $\sigma_t = 30$ ps, $\sigma_t = 5$ ps, and $\sigma_t = 1$ ps. As in Fig. 8, the orange and blue portions of each histogram represent the contributions from tumbler and non-tumbler events, respectively, while the dashed black vertical line in each panel indicates the actual value of m_1 for the

corresponding benchmark. However, we have adopted a coarser bin width of 50 GeV than we did when constructing each histogram in Fig. 8.

Perhaps the most important message of Fig. 12 is that the characteristic tumbler peak in the m_1 distribution around the true value of m_1 is evident for many of our benchmarks for $\sigma_t \lesssim 5$ ps. Indeed for BM2 and BM3, this peak is particularly striking. This once again demonstrates that an order-of-magnitude enhancement in timing resolution could yield compelling evidence of tumblers—even with an integrated luminosity that could easily be achieved at future colliders.

VII. CONCLUSIONS

In this paper, we have described a novel potential signature of new physics at colliders. This signature involves processes which we call tumblers—processes in which multiple successive decays of LLPs within the same decay chain give rise to multiple DVs within the same event. We have investigated the prospects for observing tumblers at the LHC both before and after the high-luminosity upgrade. Despite the stringent constraints that current LHC data impose on processes involving DVs, we have shown in the context of a concrete model that a significant number of tumbler events could yet be observed at the LHC. However, scenarios which give rise to a significant number of tumbler events also often give rise to a significant number of non-tumbler events which also involve multiple DVs. In order to address this issue, we have developed an event-selection procedure which permits us to discriminate efficiently between tumbler and non-tumbler events on the basis of the distinctive kinematics associated with tumbler decay chains. This procedure incorporates the timing information provided by the collider detector regarding the SM particles produced by these decay chains. As a result, the degree to which this procedure is capable of distinguishing tumbler from non-tumbler events depends crucially on the timing resolution of the detector. Interestingly, we have shown that a modest enhancement in timing precision beyond the $\sigma \approx 30$ ps timing resolution that will be provided by the CMS timing layer at the outset of the forthcoming HL-LHC upgrade could have a crucial impact on the prospects for discerning tumblers amongst possible signals of new physics involving multiple DVs. Moreover, via this same procedure, we have shown that it is also possible to reconstruct the masses and lifetimes of these LLPs. Once again, the precision to which these masses and lifetimes can be measured depends crucially on the timing uncertainty of the detector.

Several comments are in order. First, we have made a number of simplifications concerning the manner in which DVs are identified and reconstructed in our analysis. In so doing, we have accounted for the relevant uncertainties in a manner sufficient to provide a reasonable estimate of the detector capabilities necessary in order to detect a robust signature of tumblers. That said, a precise, quantitative estimate of the discovery reach for tumblers at a particular detector would require a more detailed, track-based analysis which incorporates information about the tracker geometry. Moreover, advances in detector technology may enhance the performance of particular regions of a collider detector with regard to DV reconstruction. For example, during the forthcoming high-luminosity upgrade, a High-Granularity Calorimeter (HGCal) with a timing resolution of ~ 40 ps will be installed within the endcap region of the CMS detector. This HGCal will make it possible to reconstruct DVs produced by particles whose decay products are emitted anywhere within the endcap

region of the detector with excellent precision, even at trigger level [24]. Such detector capabilities would improve the geometric acceptance for events involving DVs and therefore enhance the discovery reach for tumblers.

Second, in this paper, we have employed the mass-reconstruction procedure introduced in Sec. VI as our primary mechanism for distinguishing between tumbler and non-tumbler events. However, there may be more efficient methods of distinguishing between these two types of events. Various possibilities along these lines are under investigation [70].

Third, we have focused in this paper on the case in which the tumbler decay chains involve only three particles: χ_0 , χ_1 , and χ_2 . Indeed, this is the minimum number of χ_n needed in order to give rise to a tumbler. However, tumblers can also arise in more complicated scenarios in which the number N of χ_n particles is larger—perhaps substantially so. It is therefore interesting to consider how the tumbler phenomenology of the $N = 3$ model analyzed in this paper generalizes for larger values of N . In keeping with our established notation, we shall assume that these additional χ_n , where $n = 3, \dots, N - 1$, are all heavier than χ_2 . We shall nevertheless continue to assume that $m_\phi > m_{N-1}$. Several observations can then immediately be made.

One possibility is that the lifetimes τ_n of the additional χ_n are sufficiently short that these particles decay promptly. In this regime, tumbler events which arise as a consequence of $pp \rightarrow \phi^\dagger \phi$ production will often include additional prompt jets which can be traced back to the primary vertex. When the number of such jets is large, both triggering and the reconstruction of DVs from kinematic information becomes more challenging. Furthermore, when N becomes large, the contribution to the total event rate from processes of the form $pp \rightarrow \phi \chi_n$ and $pp \rightarrow \chi_m \chi_n$ increases simply as a result of the multiplicity of the LLPs. For sufficiently large N , the contribution from these processes to the effective cross-section $\sigma_{\text{eff}}^{(T)}$ for tumbler events—and to the effective cross sections for other classes of processes as well—can overwhelm the contribution from $pp \rightarrow \phi^\dagger \phi$.

In cases in which the τ_n for one or more of the additional states are within the DV regime, further complications arise. The reconstruction of the m_n and τ_n in this case becomes more challenging, since the tumblers themselves can involve different sequences of χ_n , even for decay chains involving only two DVs. Moreover, tumblers involving more than two DVs can also arise. Nevertheless, the methods we have developed in Sec. VI can be generalized in a straightforward manner. It is still the case, for example, that the momentum and timing information for the jets produced by a tumbler involving more than two individual decay steps is sufficient to permit the reconstruction of the m_n and τ_n of the LLPs involved in the corresponding decay chain. In particular, the reconstructed m_n distributions corresponding to the maximal tumbler decay chain—i.e., the chain involving the largest possible number of

individual displaced decay steps—will each exhibit a peak around the true value of m_n . However, the reconstructed mass distributions of *nonmaximal* such decay chains will manifest a more complicated peak structure as a result of different decay sequences involving the same number of steps. For example, a two-step decay sequence from χ_3 to χ_0 could proceed via $\chi_3 \rightarrow \chi_2 \rightarrow \chi_0$ or $\chi_3 \rightarrow \chi_1 \rightarrow \chi_0$. Following the procedures we have outlined in this paper for such two-step decays, a reconstruction of the mass of the intermediate state would then result in *two* peaks: one centered around m_1 and one centered around m_2 . Of course, given the limited spatial extent of the tracker and the fact that lighter χ_n are typically longer-lived than the heavier χ_n , decay chains involving large numbers of steps may be difficult to resolve in this manner.

Fourth, one could also consider more complicated event topologies involving LLPs which are themselves produced at DVs. Indeed, tumblers are merely the simplest example of such an event topology. More complicated event topologies in which multiple LLPs are produced at the same DV are also possible. Such possibilities would result in a proliferation of decay chains, ultimately leading to “showers” of LLPs within the collider environment. Of course, whether or not these showers are detectable as such depends on the lifetimes of the particles involved.

Fifth, in addition to considering changes in the topology of the decay chains, one might also consider changes in the properties of the individual decays themselves, such as their decay products. In this paper we have focused on models in which each decay within the tumbler produces two quarks, ultimately leading to two jets. However, it is also possible to consider models in which only a single quark is produced at each DV. In such cases, the techniques we have employed in this paper for reconstructing DVs would not be appropriate. However, as discussed above, DVs can still be reconstructed via a track-based analysis, even in such cases. Likewise, it is possible to consider models in which the SM particles produced by LLP decays include charged leptons as well as quarks and/or gluons. Methods for reconstructing DVs likewise exist for such cases.

Sixth, our primary aim in this paper has been to demonstrate that the observation of a tumbler signature is a viable possibility at the HL-LHC or other near-future colliders. Thus, while we have shown that there do exist regions of the parameter space of our example model which are consistent with current constraints, we have not undertaken a detailed analysis of exactly where the exclusion contours lie within that parameter space. Recasting tools such as MadAnalysis 5 [46,47], SModels [71], CheckMATE [72], and the computational resources associated with the RECAST framework [73,74] can assist in establishing the locations of these exclusion contours. That said, a dedicated study along these lines would be valuable, in light of the numerous subtleties involved in recasting the results of searches involving DV signatures in order to

constrain more general classes of new-physics scenarios. We leave such a study for future work.

Finally, in this paper, we have focused on the case in which both LLPs involved in our (two-step) tumbler decay chain decay within the collider tracker. One could also consider the case in which the decay of one or both of these LLPs occurs within the calorimeters or the muon chamber. Indeed, searches have been performed by the ATLAS Collaboration [37] for events involving multiple displaced decays in which one such decay occurs within the tracker and the other occurs within these outer layers of the detector. Moreover, one could also consider the case in which the lighter LLP escapes the main detector entirely and decays within an external detector designed specifically for the purpose of observing LLP decays, such as MATHUSLA [75] or FASER [76]. By incorporating information from such dedicated LLP detectors, one would potentially be able to extend an analysis of the sort we have performed in this paper across a broader range of LLP lifetimes. In fact, MATHUSLA may even be capable of detecting evidence of a tower of LLPs, as discussed in Refs. [4,77].

ACKNOWLEDGMENTS

We would like to thank Gabriel Facini and Zhen Liu for discussions. T. L. wishes to thank the EXCEL Scholars Program for Undergraduate Research at Lafayette College, which helped to facilitate this research. The research activities of K. R. D. are supported in part by the Department of Energy under Grant No. DE-FG02-13ER41976 (DE-SC0009913) and by the National Science Foundation through its employee IR/D program. The research activities of D. K. are supported in part by the Department of Energy under Grant No. DE-SC0010813. The research activities of T. L. and B. T. are supported in part by the National Science Foundation under Grant No. PHY-1720430. The research activities of B. T. are also supported in part by the National Science Foundation under Grant No. PHY-2014104. The opinions and conclusions expressed herein are those of the authors, and do not represent any funding agencies.

APPENDIX: VERTEXING PROCEDURE

We consider a pair of SM particles A and B which we assume to have been produced at the same vertex within a collider detector. We refer to the lab-frame three-momenta of these particles as $\vec{\mathbf{p}}_A$ and $\vec{\mathbf{p}}_B$, and we refer to the lab-frame coordinates at which they exit the tracker as $\vec{\mathbf{r}}_A$ and $\vec{\mathbf{r}}_B$. The trajectories of these particles lie along two lines which are described parametrically by

$$\begin{aligned}\vec{\mathbf{R}}_A(a) &\equiv \vec{\mathbf{r}}_A + a\vec{\mathbf{p}}_A, \\ \vec{\mathbf{R}}_B(b) &\equiv \vec{\mathbf{r}}_B + b\vec{\mathbf{p}}_B.\end{aligned}\tag{A1}$$

A value of a or b identifies a particular location along the corresponding line.

Given that A and B are produced at the same vertex, the lines in Eq. (A1) will intersect at the vertex location, assuming \vec{r}_A , \vec{r}_B , \vec{p}_A , and \vec{p}_B are all measured with infinite precision. However, in an actual experiment, measurement uncertainties in these quantities will typically result in the lines passing very close to each other, but not actually intersecting. We can obtain a best estimate for the intersection point by identifying the values of a and b for which the vector $\vec{D}(a, b) \equiv \vec{R}_A(a) - \vec{R}_B(b)$ is perpendicular to both lines—i.e., for which

$$\begin{aligned}\vec{D}(a, b) \cdot \vec{p}_A &= 0, \\ \vec{D}(a, b) \cdot \vec{p}_B &= 0.\end{aligned}\quad (\text{A2})$$

Solving the system of equations in Eq. (A2) for a and b , we find that

$$\begin{aligned}a &= \frac{-\vec{p}_A \cdot (\vec{r}_A - \vec{r}_B) |\vec{p}_B|^2 + \vec{p}_B \cdot (\vec{r}_A - \vec{r}_B) (\vec{p}_A \cdot \vec{p}_B)}{|\vec{p}_A|^2 |\vec{p}_B|^2 - (\vec{p}_A \cdot \vec{p}_B)^2}, \\ b &= \frac{\vec{p}_B \cdot (\vec{r}_A - \vec{r}_B) |\vec{p}_A|^2 - \vec{p}_A \cdot (\vec{r}_A - \vec{r}_B) (\vec{p}_A \cdot \vec{p}_B)}{|\vec{p}_A|^2 |\vec{p}_B|^2 - (\vec{p}_A \cdot \vec{p}_B)^2}.\end{aligned}\quad (\text{A3})$$

Evaluating $\vec{R}_A(a)$ and $\vec{R}_B(b)$ at these values of a and b and taking the midpoint between them, we obtain an estimate for the location of the corresponding vertex.

We emphasize that this vertexing procedure not only provides a way of pinpointing the location of a vertex from the measured momenta of a pair of particles produced at that vertex, but can also be used in order to assess whether or not two particles in the event were in fact produced at the same vertex. In cases in which the two particles were in fact produced at the same vertex, the magnitude of the vector $\vec{D}(a, b)$, when evaluated at the values of a and b in Eq. (A3), will be extremely small. By contrast, if the particles were not in fact produced at the same vertex, $|\vec{D}(a, b)|$, when evaluated at the corresponding values of a and b , typically will be far larger.

The processes that we have considered in this paper yield up to ten jets emanating from up to five displaced vertices when both decay chains are included. For the reasons discussed above, it is very unlikely that identifying unrelated pairs of jets as coming from the same vertex will result in small minimum values of $|D(a, b)|$. Thus by considering different pairwise combinations of jets and evaluating their minimum values of $|D(a, b)|$, it should be relatively straightforward to correctly identify those that emanate from the same vertex. We therefore expect the combinatorial background from misidentifications of jet pairs to be negligible.

-
- [1] J. Alimena *et al.*, *J. Phys. G* **47**, 090501 (2020).
 [2] O. Fischer *et al.*, *Eur. Phys. J. C* **82**, 665 (2022).
 [3] D. Acosta *et al.*, arXiv:2110.14675.
 [4] D. Curtin *et al.*, *Rep. Prog. Phys.* **82**, 116201 (2019).
 [5] P. Schwaller, D. Stolarski, and A. Weiler, *J. High Energy Phys.* **05** (2015) 059.
 [6] L. Gray and T. Tabarelli de Fatis, Report No. CERN-LHCC-2017-027, LHCC-P-009.
 [7] J. N. Butler and T. Tabarelli de Fatis (CMS Collaboration), Report No. CERN-LHCC-2019-003, CMS-TDR-020.
 [8] S. P. Martin, *Phys. Rev. D* **75**, 115005 (2007).
 [9] M. J. Strassler and K. M. Zurek, *Phys. Lett. B* **651**, 374 (2007).
 [10] T. Cohen, M. Lisanti, and H. K. Lou, *Phys. Rev. Lett.* **115**, 171804 (2015).
 [11] M. Park and M. Zhang, *Phys. Rev. D* **100**, 115009 (2019).
 [12] S. Knapen, S. Pagan Griso, M. Papucci, and D. J. Robinson, *J. High Energy Phys.* **08** (2017) 076.
 [13] R. T. D’Agnolo and M. Low, *J. High Energy Phys.* **08** (2019) 163.
 [14] K. R. Dienes, D. Kim, H. Song, S. Su, B. Thomas, and D. Yaylali, *Phys. Rev. D* **101**, 075024 (2020).
 [15] P. Giromini, F. Happacher, M. J. Kim, M. Kruse, K. Pitts, F. Ptohos, and S. Torre, arXiv:0810.5730.
 [16] M. J. Strassler, arXiv:0811.1560.
 [17] M. J. Strassler, arXiv:hep-ph/0607160.
 [18] J. E. Juknevich, *J. High Energy Phys.* **08** (2010) 121.
 [19] J. E. Juknevich, D. Melnikov, and M. J. Strassler, *J. High Energy Phys.* **07** (2009) 055.
 [20] N. Craig, A. Katz, M. Strassler, and R. Sundrum, *J. High Energy Phys.* **07** (2015) 105.
 [21] F. Lanni, L. Pontecorvo *et al.* (ATLAS Collaboration), Report No. CERN-LHCC-2020-007, ATLAS-TDR-031.
 [22] D. del Re, *J. Phys. Conf. Ser.* **587**, 012003 (2015).
 [23] J. Liu, Z. Liu, and L.-T. Wang, *Phys. Rev. Lett.* **122**, 131801 (2019).
 [24] J. Liu, Z. Liu, L.-T. Wang, and X.-P. Wang, *J. High Energy Phys.* **11** (2020) 066.
 [25] Z. Flowers, Q. Meier, C. Rogan, D. W. Kang, and S. C. Park, *J. High Energy Phys.* **03** (2020) 132.
 [26] G. Cottin, *J. High Energy Phys.* **03** (2018) 137.
 [27] K. J. Bae, M. Park, and M. Zhang, *Phys. Rev. D* **101**, 115036 (2020).
 [28] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Lett. B* **797**, 134876 (2019).

- [29] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Rev. D* **104**, 012015 (2021).
- [30] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Rev. D* **104**, 052011 (2021).
- [31] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Rev. D* **99**, 032011 (2019).
- [32] A. M. Sirunyan *et al.* (CMS Collaboration), *Phys. Rev. D* **98**, 092011 (2018).
- [33] A. M. Sirunyan *et al.* (CMS Collaboration), *Mach. Learn. Sci. Tech.* **1**, 035012 (2020).
- [34] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **97**, 052012 (2018).
- [35] M. Aaboud *et al.* (ATLAS Collaboration), *Eur. Phys. J. C* **79**, 481 (2019).
- [36] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **99**, 052005 (2019).
- [37] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. D* **101**, 052013 (2020).
- [38] A. M. Sirunyan *et al.* (CMS Collaboration), *J. High Energy Phys.* **10** (2019) 244.
- [39] A. M. Sirunyan *et al.* (CMS Collaboration), *Eur. Phys. J. C* **80**, 3 (2020).
- [40] A. M. Sirunyan *et al.* (CMS Collaboration), *J. High Energy Phys.* **05** (2018) 025.
- [41] G. Aad *et al.* (ATLAS Collaboration), *J. High Energy Phys.* **02** (2021) 143.
- [42] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **97**, 112001 (2018).
- [43] G. Aad *et al.* (ATLAS Collaboration), *Phys. Rev. D* **103**, 112006 (2021).
- [44] M. Aaboud *et al.* (ATLAS Collaboration), *J. High Energy Phys.* **01** (2018) 126.
- [45] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *J. High Energy Phys.* **07** (2014) 079.
- [46] E. Conte, B. Fuks, and G. Serret, *Comput. Phys. Commun.* **184**, 222 (2013).
- [47] J. Y. Araz, B. Fuks, M. D. Goodsell, and M. Utsch, *Eur. Phys. J. C* **82**, 597 (2022).
- [48] W. S. Cho, K. Choi, Y. G. Kim, and C. B. Park, *Phys. Rev. Lett.* **100**, 171801 (2008).
- [49] T. Han, I.-W. Kim, and J. Song, *Phys. Lett. B* **693**, 575 (2010).
- [50] K. Agashe, D. Kim, M. Toharia, and D. G. E. Walker, *Phys. Rev. D* **82**, 015007 (2010).
- [51] W. S. Cho, J. S. Gainer, D. Kim, K. T. Matchev, F. Moortgat, L. Pape, and M. Park, *J. High Energy Phys.* **08** (2014) 070.
- [52] I. Hinchliffe, F. E. Paige, M. D. Shapiro, J. Soderqvist, and W. Yao, *Phys. Rev. D* **55**, 5520 (1997).
- [53] C. G. Lester and D. J. Summers, *Phys. Lett. B* **463**, 99 (1999).
- [54] B. C. Allanach, C. G. Lester, M. A. Parker, and B. R. Webber, *J. High Energy Phys.* **09** (2000) 004.
- [55] A. Barr, C. Lester, and P. Stephens, *J. Phys. G* **29**, 2343 (2003).
- [56] D. J. Miller, P. Osland, and A. R. Raklev, *J. High Energy Phys.* **03** (2006) 034.
- [57] P. Konar, K. Kong, and K. T. Matchev, *J. High Energy Phys.* **03** (2009) 085.
- [58] M. Burns, K. T. Matchev, and M. Park, *J. High Energy Phys.* **05** (2009) 094.
- [59] K. T. Matchev, F. Moortgat, L. Pape, and M. Park, *J. High Energy Phys.* **08** (2009) 104.
- [60] K. T. Matchev and M. Park, *Phys. Rev. Lett.* **107**, 061801 (2011).
- [61] W. S. Cho, D. Kim, K. T. Matchev, and M. Park, *Phys. Rev. Lett.* **112**, 211801 (2014).
- [62] D. Kim, K. T. Matchev, and M. Park, *J. High Energy Phys.* **02** (2016) 129.
- [63] D. Debnath, J. S. Gainer, D. Kim, and K. T. Matchev, *Europhys. Lett.* **114**, 41001 (2016).
- [64] D. Debnath, J. S. Gainer, C. Kilic, D. Kim, K. T. Matchev, and Y.-P. Yang, *J. High Energy Phys.* **06** (2017) 092.
- [65] D. Debnath, J. S. Gainer, C. Kilic, D. Kim, K. T. Matchev, and Y.-P. Yang, *J. High Energy Phys.* **05** (2019) 008.
- [66] K. Agashe, R. Franceschini, and D. Kim, *Phys. Rev. D* **88**, 057701 (2013).
- [67] K. Agashe, R. Franceschini, and D. Kim, *J. High Energy Phys.* **11** (2014) 059.
- [68] G. L. Bayatian *et al.* (CMS Collaboration), Report No. CERN-LHCC-2006-001, CMS-TDR-8-1.
- [69] D. Contardo, M. Klute, J. Mans, L. Silvestris, and J. Butler, Report No. CERN-LHCC-2015-010, LHCC-P-008, CMS-TDR-15-02.
- [70] K. R. Dienes, D. Kim, T. Leininger, B. Thomas, and J. Wilhelm (to be published).
- [71] F. Ambrogio *et al.*, *Comput. Phys. Commun.* **251**, 106848 (2020).
- [72] N. Desai, F. Domingo, J. S. Kim, R. R. d. A. Bazan, K. Rolbiecki, M. Sonawane, and Z. S. Wang, *Eur. Phys. J. C* **81**, 968 (2021).
- [73] K. Cranmer and I. Yavin, *J. High Energy Phys.* **04** (2011) 038.
- [74] ATLAS Collaboration, Report No. ATL-PHYS-PUB-2020-007.
- [75] J. P. Chou, D. Curtin, and H. J. Lubatti, *Phys. Lett. B* **767**, 29 (2017).
- [76] J. L. Feng, I. Galon, F. Kling, and S. Trojanowski, *Phys. Rev. D* **97**, 035001 (2018).
- [77] D. Curtin, K. R. Dienes, and B. Thomas, *Phys. Rev. D* **98**, 115005 (2018).